

# **OPEN ARCHITECTURE COMPUTING ENVIRONMENT DESIGN GUIDANCE**

**Version 1.0 (Interim)**

**10 March 2003**

(this page intentionally left blank)

## FOREWORD

The Assistant Secretary of the Navy for Research, Development and Acquisition (ASN (RDA)) assigned the Program Executive Office for Integrated Warfare Systems (PEO IWS) with responsibility for coordinating the introduction of Open Architecture (OA) into the Navy's combat systems. Computing technology is a key part of the OA effort. Therefore, as part of the OA tasking, and based on significant research and testing, the Naval Surface Warfare Center Dahlgren Division (NSWCDD) has developed two supporting documents relevant to computing for OA. Taken together, these documents describe the technical characteristics of and requirements for computing systems in support of OA-based combat systems. This unified set of computing resources is called the Open Architecture Computing Environment (OACE). These documents provide significant insight into OACE capabilities and requirements, but they are not yet fully mature. Accordingly, they are being disseminated for formal review across the naval community.

This document, *Open Architecture Computing Environment Design Guidance, (Version 1.0)*, provides interim guidance concerning design aspects of the standards-based computing environment that is to be used in OA warfighting systems. A companion document, *Open Architecture Computing Environment Technologies and Standards, Version 1.0 (Interim)*, provides an enumeration of the standards and product selection criteria that apply to the OACE technology base. The scope of OA is intended to encompass warfighting and combat systems for ships, submarines and aircraft – including their sensor systems, weapon systems, combat direction systems and other mission critical support systems. Also included are selected C4ISR systems, the specifics of which are still under discussion. Therefore this document is intended to provide overall guidance for the design and implementation of warfighting-capable software which, when coupled with OACE, will meet warfighting mission requirements for systems across the range of deployments listed above.

This interim release of OACE information constitutes guidance that is preliminary in nature, and it will be updated according to a formal process that is being documented separately. This version of the document is being provided in order to solicit commentary from organizations outside PEO IWS. Initial review indicates that the design guidance as written, has applicability in both the combat system domain as well as the command, control and communication domains. Review of the document should include comment regarding the opportunity to provide a single guidance source for all of the aforementioned warfare domains. The guidance contained herein was made available at the beginning of the OA initiative, August 2002, for purposes of government and industry review. The formal review process for the version 1.0 final is scheduled to be complete by June 2003. An updated version of this document will be issued 30 June 2003 once that process is complete and the industry and government comments have been incorporated as appropriate. The guidance will then be updated on a periodic basis by PEO IWS according to a formal process (currently being defined) that is closely aligned with changes in the commercial market.

The interim document contains two major technical guidance sections. The first, section 3.2.2, Component Design Guidance, provides guidelines useful for constructing real-time distributed warfighting application computer programs. The second, section 3.2.3, Open Architecture Computing Environment, provides a description of the technology base to be employed in conjunction with OA warfighting systems. Based on initial industry and government feedback, it is anticipated that this document will be rewritten as two separate documents. The current section 3.2.3, Open Architecture Computing Environment will be developed as a separate document and re-titled as *Open Architecture Computing Environment Technical Capabilities*. The current descriptive content will be augmented with the addition of requirements information. This rewritten capabilities document will define computing metrics and enumerate required functionality and performance capabilities of the OACE as a computing infrastructure. The current section 3.2.2, Component Design Guidance will be recast in technical report format, and it will describe principles of distributed system design applicable to warfighting systems that fall within the scope of OA. Thus, the current section 3.2.2 should be treated as non-binding guidance. The new document will be titled *Principles of Open Architecture Application Design*.

In addition to solicitation of inputs on OA computing from non-PEO IWS organizations via this interim document, two additional avenues for solicitation of OACE input and feedback are being pursued. The first effort is a formal coordination process that gathers inputs from each PEO IWS program manager with regard to the OA effort. This program manager input process is being coordinated by Fred Riedl, NSWCCD Code B30, at [RiedlAF@nswc.navy.mil](mailto:RiedlAF@nswc.navy.mil) or 540-653-3001. As input to this process, system developers and program managers are requested by the OACE Technical Architecture (TA IPT) team to provide inputs concerning their computing requirements according to the particulars described in the next paragraph. The second effort seeks inputs from industry. The industry input effort is being managed by the TA IPT via a series of competitively awarded contracts. This latter process is scheduled to reach completion approximately end of May 2003, at which time industry inputs, as well as program manager inputs, will be incorporated into a new issue of this document and the standards document.

This interim guidance has been developed by the OA TA IPT team for Program Executive Office for Integrated Warfare Systems (PEO IWS). For further information concerning Open Architecture and the applicability of this document, contact the OA Project Officer, CAPT Thomas J. Strei, PEO IWS Code 7.0 at [StreiTJ@navsea.navy.mil](mailto:StreiTJ@navsea.navy.mil) or (202) 781-1160. With publication of this Version 1.0 (Interim) document, inputs are formally solicited from non-PEO IWS organizations. For further information concerning technical content and/or to provide recommended changes to this interim document, contact lead editor Michael W. Masters, Chairman, OA Technical Architecture Integrated Product Team (TA IPT), NSWCCD, Code B30, 17320 Dahlgren Road, Dahlgren, Virginia 22448-5100, (540) 653-1611, or via email at [MastersMW@nswc.navy.mil](mailto:MastersMW@nswc.navy.mil). Comments and recommended changes should reference a specific page or paragraph whenever possible and should provide supporting rationale describing the anticipated utility and implementation implications of the change. In addition, each responding organization should identify a single point of contact for OA related activities. Inputs should be provided no later than 15 May 2003.

# Contents

<u>Section</u>	<u>Page</u>
1 PURPOSE .....	1
1.1 Applicability .....	1
1.2 Overview .....	2
1.3 Open Architecture Goals .....	2
1.3.1 Warfighting Capability .....	2
1.3.2 Supportability.....	3
1.3.3 Open Systems and Standards .....	4
1.3.4 Managing Change with Middleware .....	5
1.3.5 Non-standard Products.....	7
1.3.6 Open Architecture .....	8
1.3.7 R&D Basis for Open Architecture .....	9
1.4 Document Overview.....	10
2 APPLICABLE DOCUMENTS .....	11
3 OPEN ARCHITECTURE GUIDANCE .....	13
3.1 Architecture Evolution .....	13
3.1.1 Open Architecture Overview .....	13
3.1.2 Open Architecture Constraints.....	15
3.1.3 Open Architecture Development Process .....	16
3.2 Open Architecture Design Guidance .....	17
3.2.1 Architecture and Design Patterns.....	19
3.2.2 Component Design Guidance .....	20
3.2.2.1 Component Design.....	21
3.2.2.2 Portability .....	28
3.2.2.3 Location transparency .....	30
3.2.2.4 Client-Server .....	31
3.2.2.5 Data Distribution.....	35
3.2.2.6 State Data Coherency .....	38
3.2.2.7 Computational Flow.....	39
3.2.2.8 Fault Tolerance.....	40
3.2.2.9 Scalability.....	45
3.2.2.10 Real-time Performance .....	50
3.2.2.11 Process, Thread and Memory Management.....	54
3.2.2.12 Data Brokers .....	56
3.2.3 Open Architecture Computing Environment .....	57
3.2.3.1 Cabling and Cabinets .....	59
3.2.3.2 Information Transfer .....	60
3.2.3.3 Computing Resources .....	62
3.2.3.4 Peripherals.....	64
3.2.3.5 Operating Systems.....	65
3.2.3.6 Adaptation Middleware.....	65
3.2.3.7 Distribution Middleware .....	66

3.2.3.8	Design Patterns and Frameworks .....	70
3.2.3.9	Dynamic Resource Management (DRM) .....	71
3.2.3.10	Instrumentation, Recording and Assessment .....	74
3.2.3.11	Failure Management .....	75
3.2.3.12	Information Assurance .....	76
3.2.3.13	Time Service .....	77
3.2.3.14	Programming Language Facilities .....	78
3.2.4	System Composition .....	79
3.2.5	Displays.....	79
3.2.5.1	Description .....	79
3.2.5.2	Guidance.....	80
3.2.6	System Test and Certification .....	80
3.2.6.1	Description .....	80
3.2.6.2	Guidance.....	81
3.2.7	Selection of Standards.....	81
3.2.7.1	Applicable Standards.....	82
3.2.7.2	Open Source Products .....	83
3.2.7.3	Non-standard products .....	83
Appendix A Glossary of Terms .....		85

# Figures

Figure 1. Open System Model .....	5
Figure 2. Managing Change with Middleware .....	6
Figure 3. Open Architecture Computing Environment.....	14
Figure 4. Functional Partitioning .....	24
Figure 5. Client-Server Model .....	33
Figure 6. Fault Tolerance Models .....	42
Figure 7. Partitioning Scalability Patterns .....	46
Figure 8. Load Sharing by Replication .....	48
Figure 9. Real-time System Comparison .....	53
Figure 10. Open Architecture Computing Environment.....	58
Figure 11. Resource Management Architecture.....	72

# Tables

Table 1. HiPer-D Architecture Concepts and Benefits .....	9
Table 2. OA Constraints .....	16
Table 3. Technical Guidance Index .....	18
Table 4. Units of Computing Performance Parameters .....	63
Table 5. Types of Standards.....	82

(this page intentionally left blank)



## **1 PURPOSE**

This Open Architecture (OA) design guidance document provides technical performance, architecture, and design guidance for the computer programs and computing infrastructure of future Naval Warfare Systems (NWS). The goal of the OA development effort is to design and build an NWS that not only continues to meet its growing and changing performance requirements well into the 21<sup>st</sup> century, but also may be rapidly and affordably maintained, refreshed, and upgraded throughout its life-cycle. Design goals also include enhanced Human Systems Integration (HSI) and optimized manning.

The purpose of this guidance document is to provide technical principles and design guidelines for redesigning and re-implementing the NWS in architecture suitable for meeting the performance, long-term maintenance, and upgrade goals described above. The guidance contained herein describes overall computing system architecture, system-wide design principles, computing equipment and support software infrastructure technologies, standards, application functional partitioning principles, computer program characteristics, and development strategy.

No change in NWS functional capabilities are anticipated in system requirement specifications such as the System Specification (A-Specification) and Prime Item Development Specifications (PIDS) as a result of the guidance contained in this document. Rather, this guidance document is intended to foster development of a robust and enduring computer program design and equipment utilization approach that will meet mission needs in the 2010 time frame and beyond.

OA will accomplish its goals by adapting and exploiting new developments in open system design principles and system architectures, as well as standards-based computing technologies from the Commercial Off-The-Shelf (COTS) marketplace. Briefly, open systems involve the use of widely accepted and available specifications, standards, products, and design practices to produce systems that are interoperable and easy to modify and extend.

The OA approach will contribute to managing, controlling, and monitoring the problems associated with technology insertion and refresh, mission capability upgrade and total ownership cost. However, it should be noted that the technical aspects of OA design are not by themselves, the total solution to all problems of system performance and cost. Additional efforts required will include process changes, new tools, revised test methods, training, etc.

OA implementation scope will be negotiated between the Navy and the prime contractor through the system development process. Where funding issues arise or requirements should be negotiated, the prime contractor will perform trade studies to facilitate cost trade-off analysis.

### **1.1 Applicability**

This document applies to the computing implementation of the functional capabilities embodied in NWSs, including but not limited to CG/DDGs, CVNs, LPDs, LHAs, LHDs, submarines, DD(X) and LCS.

## 1.2 Overview

Primary challenges facing the Navy, as documented in *Forward from the Sea* (reference (e)) and *N86 Surface Warfare Roadmap* (reference (h)), are to achieve Anti-Air Warfare (AAW) Air dominance in the face of future threat environments and to provide capabilities for Land Attack in the littoral environment. Requirements also exist to support evolution of Navy Theater Wide (NTW), an area that is of critical interest to the U.S. Navy. These challenges translate into providing surface combatants with warfighting upgrades to avoid obsolescence and to pace the threat while containing total cost of ownership.

In addition, many ship classes have major mission capabilities in ASUW, ASW, Strike, Mine Warfare and other areas. It is the intent of the OA program to consider each of these areas and to address all that are consistent with the funding and schedule objectives of the program.

## 1.3 Open Architecture Goals

In broad and general terms, architecture is defined as “the structural design of an entity.” Adding “openness” to the list of architectural characteristics implies that the “structure” of the architecture explicitly promotes interoperability, both internally and externally, as well as ease of modification and extension.

The evolution of high performance COTS, combined with continued growth of weapon system and combat system requirements, provides an opportunity to design an architecture more capable of exploiting new technologies than the federated legacy architecture that has served the Navy for well over two decades. The need for evolution toward an open architecture is motivated by both performance and supportability considerations. Commensurate with this dual set of motivating factors, the goals of the OA are as follows:

- Combat system, weapon system, command support system and HM&E capabilities that continue to pace the threat
- System design that fosters affordable development and life-cycle maintenance
- System design that reduces upgrade cycle time and time-to-deployment for new features
- Architecture that is technology refreshable despite rapid COTS obsolescence
- Improvements in NWS Human Systems Integration

Finally, system requirements may include not only capability and performance goals but also non-functional engineering goals as well (“-ilities”). In addition to traditional metrics such as reliability and survivability, OA metrics include qualitative goals such as *portability*, *scalability*, *extensibility*, and *flexibility of use*. These goals will be met, in part, by careful design and in part through use of open systems principles and standards. This document focuses primarily on the technical aspects of designing OA. In many cases, the recommended design choices and technologies are chosen with the goal of supporting this qualitative metrics as well.

### 1.3.1 Warfighting Capability

Naval mission requirements continue to grow. Newly emerging AAW threats, Area and NTW TBMD, and Land Attack support for the emerging Operational Maneuver From the Sea

(OMFTS) concept are leading to ever-increasing performance requirements for the Naval Combat Systems (NCS) and the NWS. Because of these new threats and missions, the combat system continues to evolve and change even as new destroyers continue to be constructed and existing cruisers and destroyers are modernized. The NWS computing system should be prepared to meet these expanding requirements, as well as those of other warfare areas, both now and in the future.

### 1.3.2 Supportability

In summary, the problems that are to be addressed by OA are as follows:

- Combat system requirements are increasing in scope and complexity, which dictates frequent changes and upgrades to the combat system. In turn, these lead to frequent computer equipment upgrades and computer program changes and upgrades.
- Some current combat system architectures still reflect the shared memory legacy design inherited from its military standard computing origins. For this reason, the design is relatively inflexible and difficult to change. Consequently, upgrades are both costly and time consuming.
- The implementation of the current combat system is, in many cases, specifically tailored to the technical details of the underlying networks, computers, operating systems, middleware, and other infrastructure components. Some components are non-standards-based and fluctuate significantly with respect to availability in the market place. This impacts the magnitude of computer program changes required to accommodate new COTS technologies during technology refreshes.

These factors have led to high estimated costs for life-cycle computer program maintenance, as well as equipment and support software technology refreshes. The funding required has been judged to present a significant budget barrier to continued system development and fleet support. Problems associated with high cost, long time-to-market and inflexibility of design also arises in the commercial sector. The commercial sector, in turn, has done a good job of addressing these issues in recent years.

The approach that has been developed in the commercial sector to ameliorate these limitations is often referred to as the *open systems* approach, which is elaborated in Section 1.3.3. It is likely that the open systems approach evolved first in the commercial environment because most commercial applications are considerably less stressing than military real-time applications. Most commercial enterprises do not have to contend with the additional complexities engendered by real-time mission critical systems. However, as time has passed and the computing technology has continued to improve, those solutions first adopted in the commercial sector for dealing with complexity, time-to-market, and cost containment are finally becoming feasible in real-time systems.

This practice has not necessarily been universally applied within the commercial sector; in fact, examples can be cited where commercial solutions have made matters worse. Nevertheless, many successful instances of commercial solutions are available, and it is these that form the pattern for the OA effort.

### 1.3.3 Open Systems and Standards

An open system approach has become an important aspect of system design and development in a wide variety of enterprises. This is true primarily because open systems convey certain benefits in terms of reduced life-cycle cost, reduced time-to-market, increased ability to inter-operate and cooperate with others, reduced personnel training, etc. A number of open systems definitions exist within the literature. This guidance document adopts the definition developed by the DoD Open Systems Joint Task Force (OSJTF), which operates at the level of the Office of the Secretary of Defense:

**Open system:** “A system that implements sufficient open specifications for interfaces, services, and supporting formats to enable properly engineered components to be utilized across a wide range of systems with minimal changes, to interoperate with other components on local and remote systems, and to interact with users in a style that facilitates portability.”

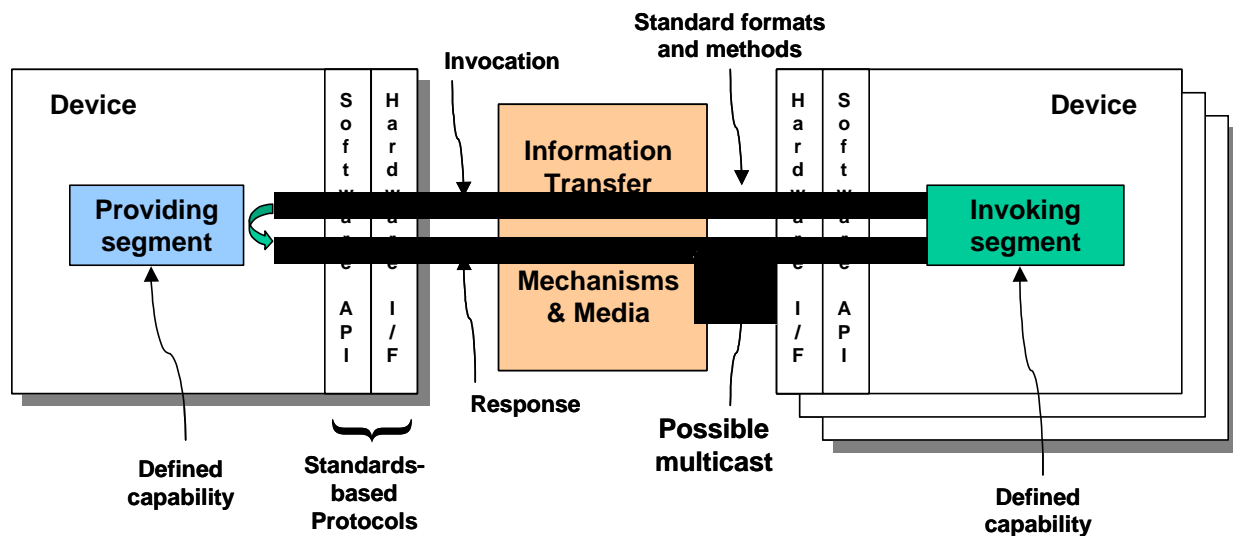
– **DoD Open Systems Joint Task Force**

Open systems – and architectures built to open system principles – possess a number of common characteristics. While not every open system possesses every possible characteristic, most open systems tend to possess most of these characteristics. Based on examination of the various open system definitions provided in Appendix A, the attributes of an open system include the following:

- Use of public, consensus-based standards
- Adoption of standard interfaces
- Adoption of standard services (defined functions)
- Use of product types supported by multiple vendors
- Selection of stable vendor with broad customer base and large market share
- Interoperability, with minimal integration
- Ease of scalability and upgradability
- Portability of application(s)
- Portability of users

A number of the points need amplification. First, while upgradability is a well-recognized characteristic of open system designs, scalability may not be so widely accepted. Scalability has a number of possible meanings, some of which are discussed later in this document. Second, open specifications are those for which a consensus-based standards organization provides control and adjudication over content and evolution. Third, the unusual attribute, “portability of users,” denotes the benefit whereby users can transition easily to a particular new system given experience with other similar systems.

The relationships among and characteristics of key components in open computing systems are shown below in Figure 1.



**Figure 1. Open System Model**

Standards play an essential role in open systems. They are the mechanism that insures that systems do not become locked into a single vendor. Note that for this principle to be effective, the standard should be widely accepted as well as formally recognized. The result of incorporating open system principles into system architecture is that such systems fulfill the following relationship:

Open Systems = Standard Interfaces + Defined Capability → Interoperability + Ease of Change

Finally, in addition to upgrading engineering practices to design Naval Warfighting Systems according to open system principles, it is necessary to define suitable metrics that provide quantitative measures of the productivity and effectiveness of the open systems approach. For example, standards compliance may be approximated by first-order estimation of the ratio of standards-based APIs to total APIs. Time and effort needed to incorporate new functions may be tracked and compared to historical data about similar activities. Numbers of interfaces, especially multi-client server interfaces, may be compared to previous implementations.

### 1.3.4 Managing Change with Middleware

. Application functions change, paced by threats and missions. Computing technology will change, driven by technological innovation and market pressures. Even standards change, albeit at a slower pace. Managing this change and isolating its effects is the key to successful OA implementation. The primary techniques for attaining isolation are use of standards, discussed in Section 1.3.3 above, and use of a category of support software known as middleware.

Middleware comes in two types - adaptation middleware and distribution middleware. Technical details of these technologies are discussed in Sections 3.2.3.6 and 3.2.3.7 respectively. A third type of support software, frameworks (see Section 3.2.3.8), provides higher level

functions such as thread management, commonly used components, etc. A fourth type, known as Resource Management (RM), provides system management functions that are not strictly middleware in nature but that, like adaptation and distribution middleware, provide vital services in the composition of large systems. Dynamic Resource Management (DRM) is discussed in Section 3.2.3.9.

Together, these four types of support software provide a synthesizing function that allows applications in a distributed system to interface with each other and with the underlying computing equipment and operating system. Figure 2 illustrates the relationship of these software services to applications and to computing equipment and operating systems.

Each of the categories in Figure 2 should, at times, undergo change. Application programs, even those designated as common services, change as missions, threats, warfighting capabilities and combat system technologies change. Computing equipment and operating systems change as computing technologies evolve and improve. Furthermore, even though middleware performs an isolation function, it too will change as the underlying computing equipment and operating systems change. In fact, the middleware's interface to the technology base *should* change if applications are to be isolated from changes in the technology base.

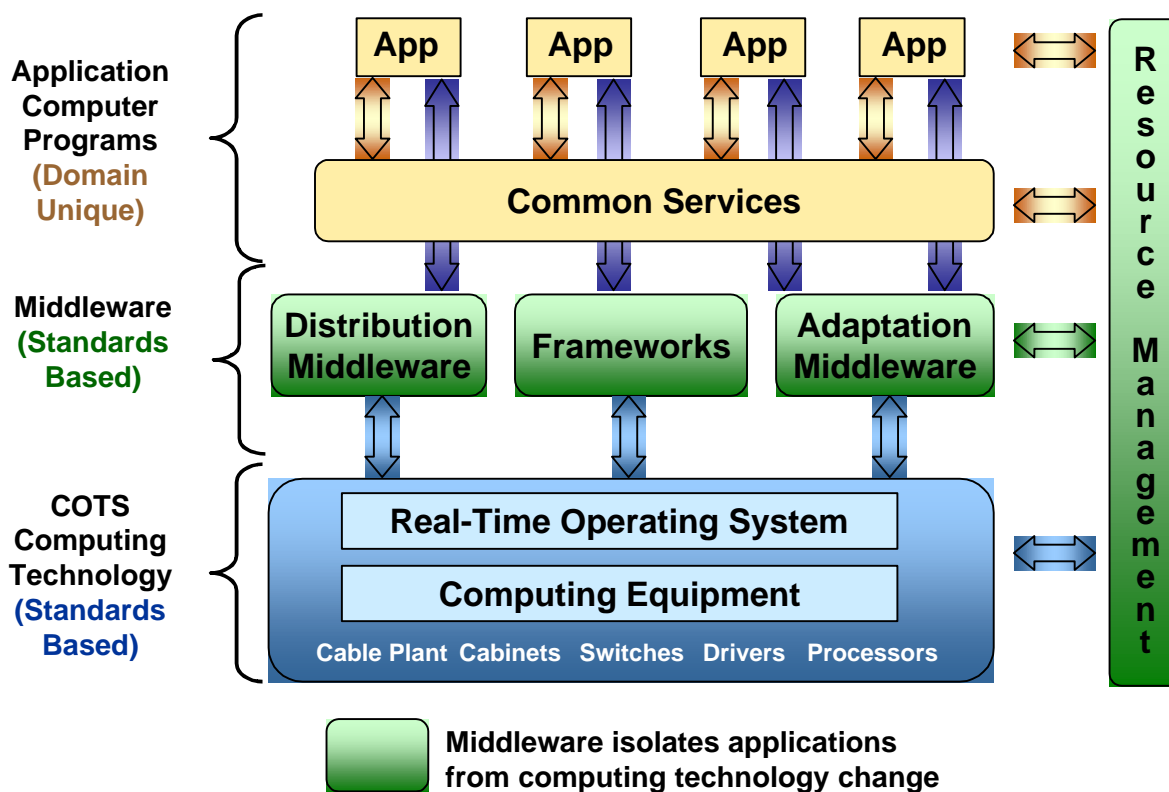


Figure 2. Managing Change with Middleware

Since middleware should change to keep pace with evolving computing technology, its implementation and support are equally as important to overall life-cycle support considerations

as are the application computer programs and the computing technology base. The same ground rules that apply to the computing technology base should apply to middleware. Middleware should conform closely to standards whenever possible. Business considerations such as market share, broad customer base, corporate stability, etc., are all valid considerations when selecting middleware products and vendors.

In fact, it is far more effective from a life-cycle supportability perspective if middleware is provided by the same commercial computing industry that also provides the remaining computing technologies. In this case, that industry and its cost accomplish middleware change and verification is shared among all customers using the technology. Otherwise, the cost of change should be borne by the developer of the system. Furthermore, since the middleware developer in this latter case is tied closely to a particular application domain, the likelihood of the middleware evolving toward a specialized and non-standard set of functions is high.

### **1.3.5 Non-standard Products**

Special discussion is appropriate for certain classes of products that support design and development of easily modified complex systems but that do not adhere to the requirement for public consensus-based, (or *de jure*) standards. Within the computing industry, a number of products exist that are not standards-based but that nonetheless have widespread market acceptance or confer certain compelling advantages to system designers, despite their lack of standardization. Often such products constitute market-based (or *de facto*) standards, which have benefits that cannot be overlooked. Potential exceptions to use of industry standards include the following categories of products:

- Products conforming to emerging standards that have not been finalized but that appear likely to do so
- Widely used products with a high degree of market acceptance and stability
- Widely used open-source products for which commercial support exists
- Non-standard products providing unique or otherwise unavailable functionality, performance, flexibility or cost benefit to the Navy

While such products frequently offer enticing advantages, it is important to note that use of these products may subsequently place system developers and acquisition managers in the position of having no competitive alternative to the products chosen. For this reason, a sound acquisition plan should require a thorough process for identification and justification of the need for non-standard products as well as management review and approval of the use of such products prior to initiation of development. When a standards-based product will serve equally, as well as a non-standards based product -- both in performance and in a business-case sense -- the standards-based product should be chosen.

### 1.3.6 Open Architecture

Section 1.3.3 introduced a number of attributes of commercial open systems. While helpful as organizing principles, these attributes do not, by themselves, constitute an architecture. Architectures are specific, albeit high level designs that satisfy not only a system's functional and performance requirements but also any relevant non-functional goals – such as open systems principles. Functional and performance requirements often determine how well or even if, the non-functional goals can be met.

Any NWS architecture – whether open or otherwise – should first and foremost meet the warfighting requirements of the warfare system and of the combat system. Fortunately, the state of computing technology has advanced to the point where satisfying both is now possible. Key NWS domain attributes and performance requirements include:

- Driven by external world events
- Wide dynamic range of processing loads
- High volume throughput of continuously refreshed data
- Low latency responsiveness to asynchronous events
- Hard real-time deadlines associated with periodics
- Computational paths traversing multiple components
- Soft real-time processing requirements
- High availability and survivability requirements
- Operator display and control requirements
- Safety and certification requirements.

Hard real-time and soft real-time characteristics are discussed in Section 3.2.2.10. To guide the design and development of an open architecture for new construction and backfit ships, the following notional definition is provided:

**Open architecture:** The high-level technical structure of the NWS as designed in accordance with the principles of open systems to achieve both real-time mission requirements and life-cycle supportability goals.

Technical characteristics of OA include:

- Distribution of processing
- Widespread use of standards-based COTS computing technologies
- Functional capabilities implemented as medium-grain components
- Use of object oriented (OO) programming within components and middleware technologies for interconnection of and interoperability among components
- Use of design mechanisms such as client-server to maximize isolation of implementation details from publicly visible services and APIs
- Portability and transparency of application components with respect to physical location and network, processor and operating system types, etc.

The corresponding goals of the OA are to provide to the NWS not only the benefits of assured technical performance, but also of reduced life-cycle cost, affordable technology refresh and reduced upgrade cycle time. Expected benefits include:



- Scalable, load invariant performance
- Enhanced information access and interoperability
- Enhanced system flexibility for accomplishment of mission and operational objectives
- Enhanced survivability and availability
- Reduced life-cycle cost and affordable COTS technology refresh
- Reduced cycle time for changes and upgrades.

With respect to survivability, an additional point should be noted. Although the current system meets its availability requirement, the OA approach combined with the dynamic aspects of shared resource management offer additional survivability benefits over those already present.

### 1.3.7 R&D Basis for Open Architecture

Much of the information contained herein is the outgrowth of a decade of shared technology and architecture evaluation and risk mitigation by the Navy, the Defense Advanced Research Projects Agency (DARPA), academia and industry. The High Performance Distributed Computing program (HiPer-D) performed system-scale prototyping directed toward the new open architectural principles being pioneered in the commercial world.

R&D efforts have focused on certain architectural concepts intended to foster life-cycle cost benefits, as well as technical performance benefits. These concepts include:

- Open architectures
- Distributed processing
- Portability
- Scalability
- Fault tolerance
- Shared resource management
- Self-instrumentation

These concepts and their benefits are summarized below in Table 1.

**Table 1. HiPer-D Architecture Concepts and Benefits**

<b>Benefit/Concept</b>	<b>Load Invariant Performance</b>	<b>Information Access</b>	<b>Mission Flexibility</b>	<b>Availability</b>	<b>Reduced Cost</b>	<b>Rapid Upgrade</b>
Open Architecture		X			X	X
Distributed	X	X	X	X		X
Portability					X	X
Scalability	X		X		X	X
Fault Tolerance	X			X		

Shared Resource Management	X		X	X		
Self-Instrumentation	X		X	X		

Open architectures provide benefits with respect to lowered development and ownership cost, rapid upgrades, and increased information accessibility. Distributed processing – as well as the ability to apply it in a scalable manner – provides the computing power needed to deliver required mission performance. Portability, one of the key aspects of open systems, reduces costs by making upgrades and technology refreshes more nearly independent of the underlying computing equipment and support software. As with any military system, fault tolerance is a vital component of system availability and survivability.

Most of the architectural characteristics listed above have sufficiently widespread antecedents to be obvious in their meanings. One possible exception is RM, the collection of services and capabilities needed to manage the computing system from a resource consumption perspective. This capability, along with the instrumentation needed to diagnose system health and to assess corrective strategies, contributes to increased performance, mission flexibility and availability. In effect, it constitutes a closed loop control system that dynamically adjusts application use of computing resources to conform to system performance and availability objectives in response to changing tactical stimuli and mission priorities.

Although not explicitly addressed in the discussion above, software reuse is an additional factor meriting strong consideration during definition of the OA implementation. Reusability may be achieved by a variety of methods, (e.g., by definition of reusable libraries and by means of application code segments designed for incorporation into multiple larger structures). Object-Oriented (OO) technology provides a very capable framework for such reuse, though OO techniques are by no means mandatory for the achievement of reuse goals.

## 1.4 Document Overview

Section 2, “Applicable Documents,” provides applicable mandates and directives that trace down to the guidance addressed in this document. Section 3, “Open Architecture Guidance,” provides the characteristics of the system that will eventually be required for its acceptance.

Within Section 3, design guidance requirements are contained in Section 3.2; design patterns are discussed in Section 3.2.1. Actual component design guidance is discussed in Section 3.2.2. The Open Architecture Computing Environment (OACE) and its technical characteristics are discussed in Section 3.2.3. Synthesis of the system from its component parts is discussed in Section 3.2.4. Display system architecture issues are discussed in Section 3.2.5. System test and certification issues are discussed in Section 3.2.6. Finally, selection of standards is discussed in Section 3.2.7.

## 2 APPLICABLE DOCUMENTS

The following documents are technical publications, directives, official correspondence, or instructions that affect the requirements put forward in this document.

- a) “Department of Defense Joint Technical Architecture”, Version 4.0, 17 July 2002.
- b) ASN (RD&A), “ASN (RD&A) commitment to Congress to completely phase out AN/UYK-43 computers by FY 99 ships”, letter, 1 December 1999.
- c) Baseline 7 Ships Characteristics Improvement Board (SCIB) Report, November 1992.
- d) Bertrand Meyer, “What to Compose,” Software Development Magazine, March 2000.
- e) Chief Naval Operations (CNO) LTR 3900 ser N865G/4U652928, 22 July 1994.
- f) Christopher Alexander, et. al., *A Pattern Language*, Oxford University Press, NY, 1977.
- g) Christopher Alexander, Sara Ishikawa, Murray Silverstein, “A Pattern Language: Towns, Buildings, Construction”. Oxford U. Press, 1977.
- h) CNO LTR ser N864/6U656004, 21 November 1995.
- i) CNO LTR ser N865G, “Aegis Combat System Baseline Plan: Annual Update”, 21 December 1996
- j) DISA, “Common Operating Environment (COE) Integration and Runtime Specification (I&RTS)”. Version 4.2, February 2002.
- k) DoDD 5000.1 of 23 October 2000, “The Defense Acquisition System”.
- l) DoDD 5000.2-R of 10 June 2001, “Mandatory Procedures for Major Defense Acquisition Programs (MDAP) and Major Automated Information Programs (MAIS) Acquisition Programs”.
- m) DoN CIO Information Technology Infrastructure Architecture (ITIA). Version 1, 16 March 1999.
- n) DoN CIO, Information Technology Systems Guidance (ITSG). Version 1, 1999.
- o) Eric Gamma, Richard Helm, Ralph Johnson, John Vlissides, “Design Patterns: Elements of Reusable Object-Oriented Software”. Addison-Wesley, 1995.
- p) Extensible Markup Language (XML) 1.0, W3C Recommendation, 6 October 2000.

- q) Guidance Document for Aegis Baseline 7 Phase 1 and II Specification Development: Information Architecture and Baseline Applicability, Version 1.0, 20 March 1998.
- r) IEEE Std 1003.1-2001. Base Definitions - Portable operating system interface (POSIX) base definitions, Issue 6, 2001.
- s) Object Management Group, “Common Object Request Broker Architecture (CORBA)”, Version 2.6.1, May 2002.
- t) PEO TSC 4890.1, “Commercial Item and Non-Developmental Item Selection, Acquisition, Integration, and Life-cycle Support Policy”,
- u) RADM Murphy Briefing, “Surface Warfare Road Map”, August 1996.
- v) SECNAVINST 5200.32A, “Acquisition Management Policies and Procedures for Computer Resources”.
- w) Secretary of Defense, “To meet future needs DOD should increase access to commercial state-of-the-art technology”, Memorandum, 29 June 1994.
- x) White Paper, “Forward From the Sea”, 19 September 1994.
- y) Interim Open Architecture Computing Environment Technologies and Standards, Version 1.0, 24 February 2003

### **3 OPEN ARCHITECTURE GUIDANCE**

The Navy intends to implement warfare systems that meet operational performance requirements and that are affordable. System affordability may be best achieved by implementing application computer programs in an open fashion on computers that are in the mainstream commercial market. The mainstream commercial marketplace typically offers a diversity of vendors, produces large quantities of products, and offers the products at competitive prices. The word *mainstream* is critical. A niche market COTS product, or one whose commercial future is uncertain, lacks most of the desirable properties stated in Section 1, even if it is driven by formal standards.

Open Architecture acts as an enabler, allowing extension of processor homogeneity (see Section 3.2.3.3) throughout the NWS, and eventually the combat system, thus allowing for the pooling of processor resources to realize the benefits associated with fully distributed computer programs. These benefits are increased scalability, improved fault tolerance and reduced logistics infrastructure.

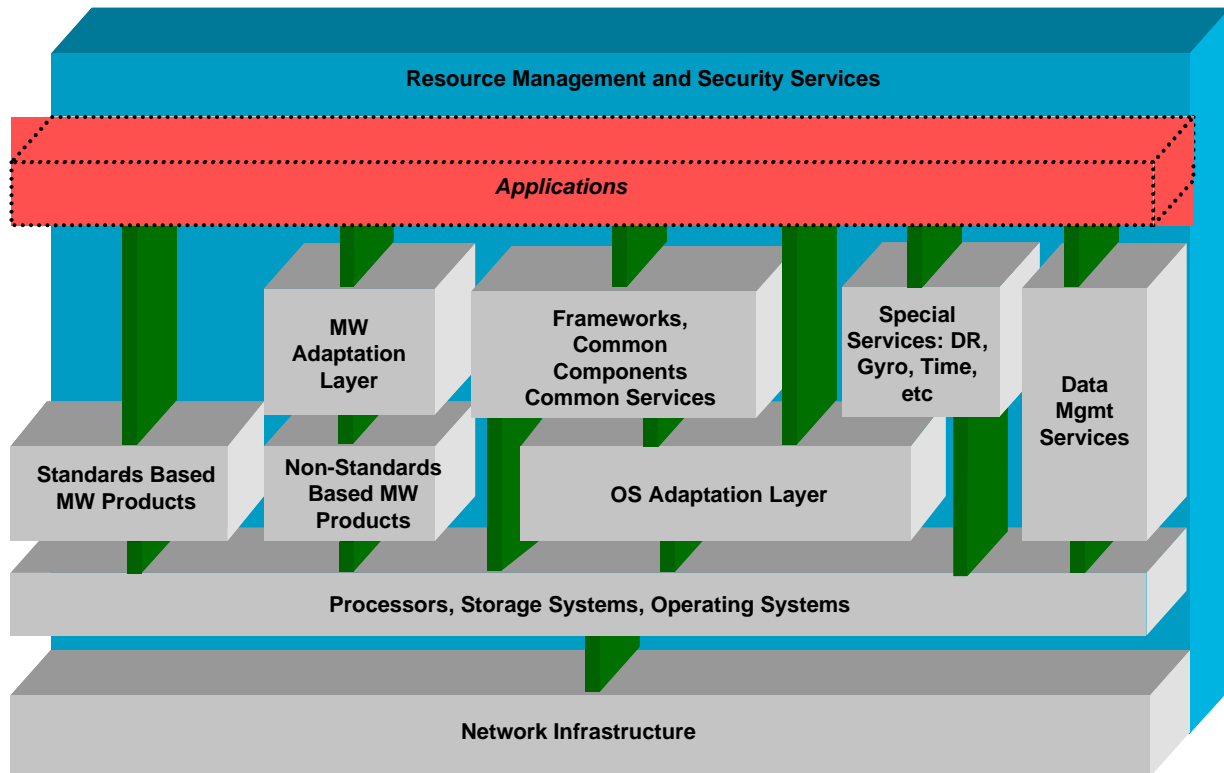
#### **3.1 Architecture Evolution**

The NWS will transition from a closed architecture to one based upon open system design principles, outlined in Section 1.3.3. Requirements for increased processing, memory, performance, and interconnectivity to support cost effective incorporation of these and future capabilities dictate this transition.

##### **3.1.1 Open Architecture Overview**

The scope of the re-designed NWS includes the underlying technology base, as well as the applications themselves. The OA design will be built on an infrastructure of technologies that include cable plant, cabinets, network components, processors, operating systems, adaptation and distribution middleware, frameworks, resource management and other services. This infrastructure is referred to in this document as the Open Architecture Computing Environment (OACE). OACE (as described in Section 3.2.3) will allow the Navy to introduce and change out commercial technology to maximize affordability and performance goals. Figure 3 illustrates the relationship of various technology components and the applications they support.

The starting point for OA development process is a partitioning of NWS functionality that will create a new component-based logical architecture no longer limited by design constraints arising from the previous generation of computing technology, as reflected in the current element-based architecture. The selection of partitioning boundaries will be accomplished by examining the entire NWS functionality set during the design stage. However, implementation may be accomplished in phases consistent with funding availability. Due to constraints imposed by the legacy architecture, selection of implementation phases should necessarily take current element boundaries into consideration.



**Figure 3. Open Architecture Computing Environment**

The logical architecture will be defined so as to incorporate all specified NWS capabilities and requirements. These capabilities are defined by the appropriate Operational Requirements Documents (ORD) and by the Combat System Specification Requirements.

It should be noted that the nature of OA partitioning is likely to be such that new capabilities not previously affordable may become available on a shorter and lower cost cycle than in the past. This is, in substantial measure, true because information within an open NWS implementation will be much more readily available than in current baselines. The use of networks, client-server implementations, general purpose computing equipment and services, etc. allows convenient access to the information and computing resources needed to incorporate new functionality. This, in turn, has the potential to provide a computational base for increased automation in support of manning reduction.

Furthermore, the use of a “pool-of-computers” architecture for the resource base, combined with the application of DRM, has the potential to change the approach to shipboard computer system maintenance. It may be possible to substantially reduce shipboard computer maintenance by capitalizing on the fact that application components are not bound to computer locality but instead are free to migrate to available processors under RM control. Reduction of shipboard maintenance should lead to reduced manning.

The logical architecture should be as technology-neutral as possible. The purpose of this approach is to ensure that long-term supportability goals are met. Only after the logical architecture is defined should an effort be made to define areas where code reuse is possible.

Implementation will utilize best commercial design and development practices and state-of-the-practice technologies consistent with NWS mission requirements.

Mission critical enclosures will house commercial-based advanced processors, network, and conversion equipment. The legacy equipment footprint will be retained for OA. That is, the number and placement of enclosures, as well as inter-compartment cabling, will remain fixed. The NWS computing capability should be capable of uninterrupted operation under battle conditions that include but are not limited to shock, fire, and salt water spray. This should be accomplished in a manner that is survivable, maintainable, and affordable.

However, a “pool-of-computers” approach will, in general, allow any computer program to run on any available processor – subject to limitations imposed by high performance real-time requirements. Fault tolerance and survivability will be based on an N+M redundancy scheme, where N is the number of processors required to perform mission computing tasks and M is the number of spare processors.

NWS communications will be based on LAN technologies. Limited Navy Tactical Data System (NTDS) legacy interfaces will continue to be supported. However, where feasible, NTDS legacy interfaces should transitioned to the OA LAN-based architecture.

Component boundaries and interaction requirements will be established in conjunction with and consideration of timing and performance characteristics. The implementation and deployment architecture will be established by finding the best technology that satisfies overall design objectives. The implementation architecture deals with technology choices and grouping of components to minimize delays or latency based on the technology constraints. Subsequent baselines will use the same architectural framework to select new technologies. This allows the components to be reallocated as technology constraint change.

### **3.1.2 Open Architecture Constraints**

OA represents an opportunity to reinvigorate the performance and maintainability of NWS with a new functional allocation, a computing architecture and mainstream standards-based computing technologies. Likewise, the investment in physical design and construction of the surface combatants should be considered. Thus, while it is possible to change computer programs, and even network and processor components, without impacting ship construction, it is considerably more difficult to change cabling and cabinetry. For this reason, once the best technical design is identified, cost-benefit tradeoffs will evaluate any proposed physical design changes that may impact physical ship construction and modernization costs and schedules.

With the above caveat in mind, Table 2 lists the initial OA design constraints. In the table, the term “mainstream” captures an important concept that contributes to achieving OA objectives. Generally, it connotes the idea of product lines with large market shares, extensive user base, readily available support and vendor stability. Standards bodies embraced for the OA initiative include but are not limited to Telecommunications Industry Association (TIA), Internet Engineering Task Force (IETF), Object Management Group (OMG) and the IEEE’s Portable Operating System Interface (POSIX) standard group.

**Table 2. OA Constraints**

<b>System Domain</b>	<b>Constraints</b>
Cable Plant	Use TIA standards-based COTS cabling technology when possible
Enclosures	Use enclosures capable of providing non-mil-spec environment for mainstream COTS components
Consoles	Use consoles consistent with Navy acquisition process
Networks	Use mainstream IETF standards-based COTS network technology
Computers	Use mainstream COTS processors and boards
Operating Systems	Use mainstream POSIX standards conformant operating systems
Middleware	Use mainstream OMG standards based middleware
Resource Mgt.	Implement NWS-wide static and dynamic capabilities
Common Services	Implement common support services as system server applications (e.g., time services, data extraction, navigation data distribution, etc.)
Application Design	Redesign tactical applications as components in accordance with open systems concepts outlined in this guidance document

### **3.1.3 Open Architecture Development Process**

Effective processes methodologies and tools are critical to the success of large scale engineering projects. Key characteristics of the incumbent process include a strong emphasis on firm definition of requirements prior to design and coding as well as a formal design review process involving Navy buy-off at every stage of the development. When required warfighting changes were extensive in scope, EDMs usually preceded production development as a means of risk mitigation – the latest of which will be an OA EDM.

The Object Oriented (OO) design method did not originate in the real-time community, but it has matured to the point where the performance characteristics of OO computer programs are, in many cases, more than adequate to meet real-time requirements. The OO approach not only supports critical OA objectives, such as affordability and time-to-market, but also represents state-of-the-practices among trained engineers within the pool of likely OA implementers.

Given this situation, the time has come to update processes, methodologies and tools to the current generation as a part of OA development.

The new spiral approach is far more inclusive and graphically oriented than previous approaches, due in no small measure to the degree of automated and collaborative design that is



now possible. Unified Modeling Language (UML) and OO allow engineers to express and capture requirements and design in a readily transportable and interchangeable form. Web-based Integrated Data Environments (IDEs) provide immediate access to project artifacts. When coupled with a management philosophy that promotes early and broad access to intermediate products by Integrated Product Teams (IPT) via the IDE, this approach allows maximum visibility to implementers and customers alike.

Notwithstanding the advantages of new tools and methods, perhaps their most important advantage is the ability they convey to overlap requirements setting, system engineering, design and implementation successfully. In the past, substantial parallelism across these stages was considered off process and was often not formally allowed. This resulted in delaying program direction changes in response to problems uncovered during development, with consequent schedule and cost disruption when such problems were inevitably encountered. However, where tools permit and where traditional engineering rigor and discipline are sustained, such overlap provides a useful mechanism for identifying and addressing problems early in the development cycle.

Having stated the preceding benefits of new tools, understanding how the process should be applied is important if these benefits are to be achieved. Two criteria for selection of candidates for early definition and prototype implementation include high-risk capabilities and foundation functionality. Each has its merits. In the case of high-risk functions, the benefit is obvious. The more time and effort devoted to resolving risks, the more likely is overall project success. Likewise, foundation functionality is capability upon which subsequent system operation should depend. Examples include technology-based infrastructure components, such as networks, operating systems, and middleware and resource management; and core system functionality, such as track management and common tactical picture.

The traditional structured approach to system engineering and design might reasonably be characterized as “top-down” and “breadth-first”, with a major focus on a sequential customer/developer review and approval process exercised on a step-by-step basis. By contrast, the new approach is simultaneously top-down but with “depth-first”, or even “bottom-up” excursions in critical risk and foundation areas. The new process also includes a major focus on system architecture. In summary, the new methods and tools enable a more fine-grained and risk-aware development process without implying or demanding compromise of the disciplined top-down approach that is integral to successful large-scale system engineering endeavors.

### **3.2 Open Architecture Design Guidance**

The technical guidance points discussed herein constitute the critical technical characteristics needed to develop the OA; therefore, they are the primary substances of this guidance document. However, they may differ in certain aspects from a more general description of open systems since these attributes might be applicable to a broad base of non-real-time and non-mission critical applications. These differences reflect the more stringent requirements of real-time mission critical systems.

The level of detail provided in the following sections is, in general, below that associated with requirements. The intent of the guidance is to characterize design techniques that have

proven successful in other systems and prototypes comparable to NWS. Unless otherwise noted, guidance alternatives are not necessarily all-inclusive. Thus, additional implementation possibilities may exist.

Section 3.2.1 introduces the concept of design patterns; i.e., ways of building software that are useful in many different forms across a wide variety of applications. Technical guidance for OA is divided into two categories: the attributes of application computer programs and the attributes of the underlying technology base. Section 3.2.2 provides a detailed discussion of component design characteristics and constraints. Section 3.2.3 covers the characteristics of the underlying technology base. Table 3 provides an index to each discrete guidance item in the document.

**Table 3. Technical Guidance Index**

<b>Category</b>	<b>Paragraph</b>	<b>Page</b>
Component design	3.2.2.1	21
Portability	3.2.2.2	28
Location transparency	3.2.2.3	30
Client server	3.2.2.4	31
Data distribution	3.2.2.5	35
State data coherency	3.2.2.6	38
Computational flow	3.2.2.7	39
Fault tolerance	3.2.2.8	40
Scalability	3.2.2.9	45
Real-time performance	3.2.2.10	50
Process, thread & memory management	3.2.2.11	54
Data Brokers	3.2.2.12	60
Cabling and Cabinets	3.2.3.1	59
Information Transfer	3.2.3.2	60
Computing Resources	3.2.3.3	62
Peripherals	3.2.3.4	64
Operating Systems	3.2.3.5	65
Adaptation Middleware	3.2.3.6	65
Distribution Middleware	3.2.3.7	66
Frameworks	3.2.3.8	70
Dynamic Resource Management	3.2.3.9	71
Instrumentation	3.2.3.10	74
Failure Management	3.2.3.11	75
Information Assurance	3.2.3.12	76
Time Service	3.2.3.13	77
Programming Language Facilities	3.2.3.14	78
Displays	3.2.5	79
System Test and Certification	3.2.6	80
Selection of Standards	3.2.7	81
Open Source Products	3.2.7.2	83

### 3.2.1 Architecture and Design Patterns

Patterns are widely useful reusable rules and methods for building components and systems, including open systems. Architecture patterns constitute high-level structures appropriate to design of major segments of software, such as computer programs and components. They may foster key open system properties such as code reuse, flexibility of design, enhanced interoperability, etc. Design patterns are a more detailed and discrete templates. The following definition of patterns comes from Alexander (1997). Although it deals with patterns in buildings and towns, nonetheless, its generality makes it highly applicable to computer systems as well.

“Each pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice.”

Extending this definition, designers Gamma, Helm, Johnson and Vlissides state the following in their book, *Design Patterns: Elements of Reusable Object-Oriented Software*: “[A]t the core. . . of patterns is a solution to a problem in a context.” They list the following four characteristics as the essential elements of a pattern:

- The pattern name is a handle we can use to describe a design problem
- The problem describes when to apply the pattern, and the pattern explains the problem and its context
- The solution describes the elements that make up the design, their relationships, responsibilities, and collaborations
- The consequences are the results and trade-off of applying the pattern

Certain important design patterns recur in open systems. The patterns listed below, as well as the tools required in their construction (largely middleware) have shown that they have value in the real-time weapon system problem domain. They will be discussed in more detail in subsequent sections. In many cases, application components may exhibit aspects of more than one pattern. The key patterns include:

- Client-server patterns
- Distribution patterns for continuously refreshed data
- State data coherency maintaining patterns
- Computational data flow patterns
- Replication patterns for fault tolerance
- Peer-client replication pattern for load-sharing and scalability
- Inter-component adaptation patterns (e.g., brokers, translation layers)

Patterns describe a particular way of solving a problem. The pattern may be a high-level abstraction of a solution, such as client-server, or it may be a very detailed recipe for solving certain classes of problems. In any case, the pattern is definitely a solution and not just a description of the desired result or the environmental constraints.

### 3.2.2 Component Design Guidance

In OA, application computer programs should be implemented as loosely coupled software components. Each component is a separately compiled and linked executable computer program, or process in POSIX operating system terminology. In a full OA implementation, a large-scale infrastructure of processors, networks, and services supports the distribution and execution of application components such that both scalable performance and system fault tolerance are achieved.

In order to facilitate the evolution of advanced computing in naval combatants over the lifecycle and to achieve the benefits of distributed computing the tactical computer program technical attributes described by this guidance document should be implemented, to the greatest extent possible, by all NWS elements. Whenever technically and programmatically feasible, the tactical computer program technical attributes will also be extended to ACS elements.

Attributes of the OA application computer programs are listed below and described in the subsections that follow.

- Functionally distinct self-contained applications or components, usually modest in size
- Components loosely coupled in space and time with other components
- Applications built for portability and location transparent allocation and operation
- Client-server design patterns with push and/or pull interfaces
- Mechanisms for distribution of continuously refreshed data
- Maintenance of state data coherency in components with composite state
- Computational flow and multi-component paths
- Multiple fault tolerant design patterns (active, passive and hybrids)
- Scalability via load-sharing peer-client design pattern
- Quality of Service (QoS) mechanisms for support of both hard real-time and soft real-time requirements, possibly collocated
- Multi-threaded applications (I/O, periodics, etc.) with threads not hard mapped to Central Processing Units (CPUs)
- Pre-allocation of computing resources (memory, processes, threads, etc.) at application initialization
- Mechanisms for preservation of data integrity (correctness and resistance to disruption) across threads, processes, computers and networks
- Asynchronous processes with timed/fixed packet I/O buffering
- Push distribution (forward caching) of track data to consumers
- Data broker/adaptor design patterns for legacy capture
- Any-operator, any-console display architecture

The attributes of the technology base needed to support these application component attributes are described in Section 3.2.3. For the OA, the following global guidance should be incorporated into each tactical component, as applicable: all NWS functions should be implemented as distributed components designed specifically for the OA computing environment. Non-NWS elements are encouraged to migrate to OACE; however, this is not required. Individual attributes are discussed in Sections 3.2.2.1 through 3.2.2.12.

### 3.2.2.1 Component Design

This section provides rationale and requirements for partitioning NWS functional capability into medium-grain, functionally homogeneous, loosely coupled code units, or components that are suitable for distributed processing.

#### 3.2.2.1.1 Description

A key concept in structuring application computer programs for incorporation into OA is the advantageous partitioning of system functionality into code units, or *components*, that are easy to change and to interoperate with over the life-cycle of the system. In this characterization, a component is a physical replaceable software entity or element that performs defined functions and encapsulates data. A component is a distinct, separable and uniquely identified constituent part of a system, with well-defined boundaries, capabilities and interface to other system components. Note that, conceptually, components may be either hardware or software. However, in this context, only software components are intended. Both processes and libraries may qualify as components.

**3.2.2.1.1.1 Utility of Components.** Components provide access to encapsulated functions and data via well-defined interfaces, and they exhibit predictable behavior. Components are usable by multiple heterogeneous clients without modification. Usually, process-based components are defined along boundaries that facilitate loosely coupled communications with other components. Moreover, the scope of a component should usually represent a lowest common denominator of capabilities, such that little, if any, further decomposition can be performed on the functionality while still preserving the loosely coupled character of the component.

Components are used to compose larger aggregations. The entire set of components aggregated for an overall purpose constitutes a *system*. However, smaller collections of components often arise as well. These smaller collections may have a degree of functional unity and utility that merit separate identification and tracking. Such collections of components are referred to as *subsystems*. One example would be a data extraction and recording subsystem, which itself may be composed of a number of separable components. Just as components may be part of multiple systems, so subsystems may also be constituent parts of multiple systems.

If a software element is understood generally as any collection of software built for a specified purpose, the notion of components introduces constraints on the choice of boundaries and structure – constraints that are intended to foster ease of composition, reuse and other supportability metrics. Also central to the notion of a component is the concept of *durability*. The durability metric suggests that functionality, Application Programming Interfaces (APIs) and inter-component interfaces should exhibit some considerable degree of stability over time and through repeated functional enhancements to other system components.

When properly designed, components have a number of key advantages. They facilitate construction of reusable code units that can be inserted easily into a wide variety of systems without requiring knowledge of the internal implementation of the component. Furthermore, the

stable component interface facilitates the evolution of a system through specialized, extended elaborated or optimized versions of the component.

**3.2.2.1.1.2 Composing Systems with Components.** According to Bertrand Meyer (2000), “the two basic conditions for a software element to be considered a component are that it be:

- Usable by other software elements. This excludes a program in the traditional sense that is meant to be used by humans or non-software triggers – unless it has been componentized, meaning precisely adapted for use by other software.
- Usable by software elements, (whose) authors are unknown to the component's authors. This excludes the case of routines, classes and other software elements used by other parts of the same software. A component should be of interest to a broad range of ‘clients’ not directly connected to the original authors.”

The preceding points have direct applicability to the OA initiative in that OA has as a major objective getting to a point where the Navy “buys once” and reuses components across multiple ship classes. In any case, according to the author, “These requirements, modest as they may seem, immediately lead to several others.” Among them are the following:

- A component should include a specification of all its dependencies: hardware and software platform, versions and other components. Otherwise new clients will not be able to make good use of the component without going back to the original author.
- For the same reason, a component should provide a precise specification of the functionalities that it offers.
- The component should be usable on the sole basis of that specification, without access to non-interface information (such as the source code even if it is available). This leads in particular to the information hiding requirements.
- Components should be composable with other components, since a single component is not very exciting and certainly does not justify talking about component-based development. *In practice, this means that a good component will usually be part of a more general component framework with a clear overall architecture, style and standard design patterns.* [emphasis added]
- The process of integrating a component into systems should be fast and smooth.

With regard to the composability characteristic above, it should be noted that integration of a component that supplies a new service requires creation of new clients and will always take longer than integration of a new client or a server that takes over an existing interface. In any case, components are usually loosely coupled to other components (i.e., communicate via message passing); however, they may be tightly coupled internally if system performance needs so require. Based on the preceding discussion, component requirements include the following:

- A well-defined behavior that is reflected through services offered the component
- Encapsulation of the data that is managed by the component
- A well-defined and durable interface through which services are obtained.

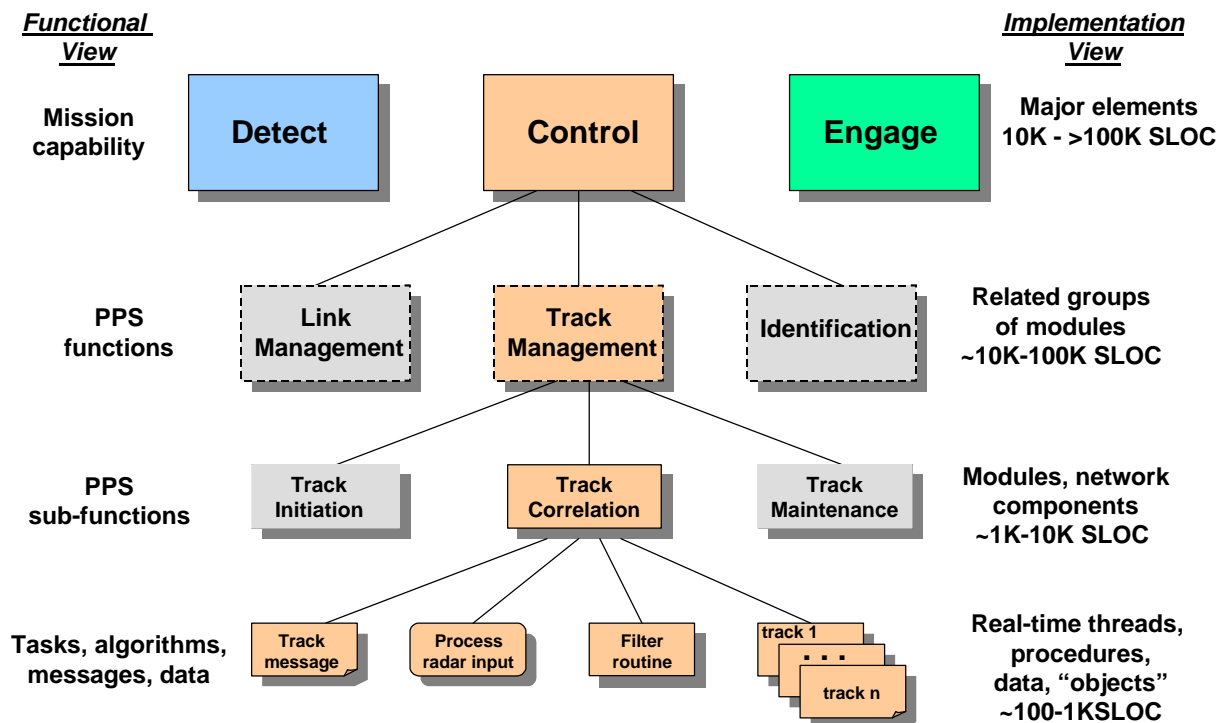
In OO design technology terminology, active components may be loosely thought of as large-scale executable objects designed to be easily replaceable as a unit in the context of a system. Especially beneficial is the situation when the interface is amenable to formal characterization, (e.g., by Object Management Group [OMG] modeling language specification or the UML). However, while the term, object is often applied to such components, use of OO technology is not, in the strictest sense, a requirement for effective system partitioning.

**3.2.2.1.1.3 System Partitioning for Components.** Partitioning NWS functionality into components departs from earlier system development efforts. It tends to produce smaller, more functionally homogeneous and loosely coupled program units rather than a few large non-homogenous shared-memory computer programs. These programs tend to be medium-grain rather than fine grain, as might be found in data flow decomposition.

Figure 4 provides a notional representation of this distinction. This diagram is consistent with both a functional and an OO view of system design. Two points should be considered: First, a pure OO view is neither necessary nor desirable in defining components. In fact, components will likely take on characteristics of each. It is more likely that components will be program units with OO interfaces than purely OO. However, Figure 4 does lend support to the notion of OO development, particularly at the finest level of granularity.

The terms “functionally homogeneous” and “loosely coupled” should be understood in general rather than precise terms. In this context, these terms imply a number of things about the structure and composition of the body of code that constitutes the component, as follows:

- The boundaries of components result in clearly distinct entities separated by considerable functional distinctiveness and distance in a requirements sense
- Components may rely on shared services provided by other components; in the example above, both data link processing and NWS control need access to track information; this in turn suggests yet another possible component boundary, in this case, one that provides service to multiple track-consuming components.
- The operations and/or services provided by the component that other components are logically closely interrelated and mutually supportive.



**Figure 4. Functional Partitioning**

Component boundaries are chosen to maximize internal binding strength while minimizing external binding strength. In this case binding strength is related to the frequency, volume and timeliness requirements of data access among code segments. Based upon component boundaries selected in accord with the binding strength criterion:

- Intra-component data structures and operations are so closely related that direct visibility and tight coupling of processing and data access (e.g., via shared memory) is appropriate internally; at the same time, loose coupling (message passing and distributed object invocation) to all other components is practical in a performance sense.
- Encapsulation, scoping, and visibility of externally referenced data structures are chosen so that an API-based non-shared memory, non-direct invocation style of system construction may be maintained without rendering component performance unsuitable and without exposing non-essential internal data to external visibility and access.
- When senescence and bandwidth characteristics permit, tightly coupled (e.g., shared memory) access to data needed by multiple code segments is replaced by “forward caching” of data to consumers. This is done via an explicit multicast-like data distribution mechanism (i.e. distribution of data by a server-like entity in advance of its need by consuming clients). By this forward caching mechanism, data references that were formerly accomplished by explicit access to memory shared are converted to local references within the component using the forward cached data.
- Size of the component code base is a few hundred to a few thousand lines of source code (i.e., “medium grain”). Exceptions to the size guideline involve large, tightly coupled



applications for which loosely coupled partition boundaries are not available.

By the conventions used herein, components may be either active or passive. Thus, components may be processes (active), or they may be libraries (passive). Design considerations for active components include use of priorities, control of priority inversion, thread safe synchronization and data access, and replication. Design considerations for passive components include re-entrance. In both cases, error handling should be considered.

The most common view of an active component is that a component constitutes a separate process. However, a capability called collocation is available within some distributed object-computing models. With collocation, objects may either be located on the same machine and within the same address space as other code segments, or they may be located remotely, accessible only via physical distributed computing mechanisms. However, code built using collocation mechanisms are less portable than non-collocation code, therefore collocation is generally not recommended unless no other mechanism will meet application requirements.

Component technology such as Java Beans and the CORBA Component Model (CCM) may provide additional support for building reusable components and extensive services. However, Java Beans are only useful where Java is selected as the implementation language. The CCM has the advantage of providing component support in a multi-language environment.

#### **3.2.2.1.2 Guidance**

The OA tactical computer program architecture will consist of a set of distributed components whose collective interactions provide the complete specified functionality of a required operational capability. Distributed components have available the underlying technology base specified throughout Section 3.2.3. Each member of a set of fully distributed software components should be capable of execution within a common unit of hardware. Exceptions to this rule may occur in a large embedded real-time system. Thus, some components may exhibit dependence on unique input/output or special computational resources, such as a group of closely coupled processors. Except as required because of limitations involving unique resource and/or connectivity requirements, components should make use of general purpose OA computing and connectivity resources. Distributed components should exhibit the following characteristics:

1. Components should preserve NWS functionality and real-time performance.
2. Component functionality and interface definitions should be verified and documented via explicit and comprehensive use case analysis, as exemplified in the UML and associated design tools. Analysis should include real-time performance, latency and timing considerations using real-time UML extensions as appropriate.
3. Components should fall into one of several categories: (1) compiled, linked, stand-alone computer programs, (2) library elements, which can be loaded, and instantiated by such linked components. (3) Nested components wherein simple components are incorporated into a more complex compound components, (4) reusable services and components, (5) COTS packages, etc. Compound components involving nesting may be created in a variety of ways, including linking in of library services, instantiation of simple classes with inheritance of simple class properties, etc.
4. Component boundaries should maximize intra-component cohesion, coupling, and

inherent parallelism. In order to leverage the unit of computing concept (see Section 3.2.3.3), and within the constraints of required performance and internal cohesion requirements, encapsulated component functionality should be a small, manageable subset of the overall functionality of the particular tactical element from which it derives. As a rough guideline, simple components should be “medium grain” (i.e., sized on the order of a few hundred to a few thousand lines of source code). Where necessary due to the monolithic nature of certain processing algorithms, simple components may also be larger.

5. Partitioning of tactical functionality should be based upon an OO and/or structured (control flow) decomposition method. This not only minimizes data flow between components, but this method also limits ripple effects into other components of internal component changes. Thus, it supports reuse and enhances maintainability.
6. Components should be as self-contained and functionally homogeneous as possible and should encapsulate closely related and logically coherent functionality and data.
7. Components should be designed for self-contained and independent operation under control of the host operating system. Components should be multithreaded where necessary to preserve time and event scheduling properties.
8. Components should be designed to have low cohesion (i.e., a loosely coupled, message passing or distributed objects interface) with other components in the system.
9. Components should be designed such that their reliance on data from other software components in the system is minimized. When data sharing among components is required, techniques such as or distributed objects or message passing should be utilized to the greatest extent possible in order to retain a loosely coupled system.
10. Component interfaces should be well-defined and durable as to content, function and methods of implementation. Interface definition should be made at the logical level rather than the physical level using standards-based interface definition languages and automatic interface code generation tools. The various technologies that contribute to developing inter-component interfaces are discussed in Section 3.2.3.
11. Component interfaces should be defined in an appropriate interface specification document or via an Interface Design Language (IDL). The API for this interface should be defined in an acceptable language or meta-language (e.g., Ada package specification, C or C++ header file or CORBA IDL).
12. Where efficiency considerations permit, standards-based self-defining interface methods (e.g., Extensible Markup Language (XML)) may be employed. For example, XML might be used to provide a variable format status data to an operator for display.
13. Components should not use explicitly shared memory techniques to communicate data between components; however, they may use shared variables among multiple internal threads so long as proper data protection primitives are employed. This should not be construed to forbid use of shared memory between an application and the OS or between application and co-resident support functions such as timekeeping, data extraction, etc. When shared memory use is absolutely unavoidable due to performance considerations, standard methods should be employed.
14. Explicit control flows between components should be minimized, especially where mechanisms for information accessibility (e.g., publish-subscribe) are available that do not require symmetric (i.e., two-way) binding of component logical identities. Such accessibility mechanisms promote server interfaces that may be shared by multiple

clients rather than a pair-wise interface approach. In the pair-wise approach, a unique interface is built between every two components that should communicate, thus increasing cost and complexity. Such non-symmetrical information access mechanisms also promote component decoupling since the server need not know anything about the clients.

15. All OA components with time critical performance and recovery requirements should be replicated for fault tolerance. Fault tolerance mechanisms are discussed in Section 3.2.2.8. Components with less critical recovery time requirements may be designed so that they can be restarted under RM control (see Section 3.2.3.8).
16. All components for which potential future performance requirements exceed single processor computing capacity or network bandwidth should be replicated for scalability. Scalable load sharing is discussed in Section 3.2.2.9.
17. Components subject to replication (i.e., for fault tolerance or scalability) should be designed so that several active instantiations (or replicas) of the component can execute simultaneously in physically different processors/locations.
18. When multiple active instantiations of a component are configured for fault tolerance, the instantiations should provide a means for coordination of state information between the replicas and for deciding which copies fulfill defined roles such as primary, backup, etc.
19. When multiple active instantiations of a component are configured for load sharing, the RM capability (see Section 3.2.3.8) should redistribute load and instantiate or remove active replicas for purposes of achieving system load balancing, as well as fault tolerance. Therefore, load balancing application components should be designed in such a manner to facilitate RM execution of the requisite control action. Application replicas should implement and coordinate the load balancing strategy. RM will add and/or remove replicas according to that strategy. Further investigation is required to define specific mechanism.
20. When multiple components perform functionally related operations (e.g. similar processing functions for different sensors, weapons or warfare areas), an investigation should be conducted during component design to assess reuse possibilities. Alternatives to be considered include taking advantage of reuse possibilities inherent in the similarity of the components (both in the form of design pattern reuse as well as code reuse) as well as implementing the component as a common support service.
21. When multiple mature implementations of the CORBA Component model are available, then use of CCM should be considered. For Java implementations, the Java Beans model should be considered.
22. Real-Time QoS metrics should be defined in terms of component performance requirements and multi-component execution paths determined by domain analysis. Real-Time QoS metrics and thresholds should be identified and quantified for each distributed tactical function hosted by OA. Real-Time QoS metrics should be utilized by the RM capability to dynamically balance the loading in the system as described in Section 3.2.3.8. Further investigation is required to define specific mechanisms.
23. Components should be robust with respect to exception conditions. Where appropriate, exceptions should be detected and handled internally to each component. When necessary, exceptions should be defined within the component's interface and conveyed to invoking code segments in other components.

### 3.2.2.2 Portability

This section provides rationale and guidance for designing OA components to ensure portability of source code across networks, computers, operating systems, middleware, and other technology base components of the OA computing system.

#### 3.2.2.2.1 Description

A major characteristic of today's computing industry is the fact that the technology base is changing and evolving rapidly. This potentially provides great benefit in terms of a steadily improving price/performance ratio. However, systems that are not designed to accommodate this rapid change of the technology base can be quite costly to maintain over the life-cycle. Application source code not designed for portability should be modified, sometimes extensively, when it is ported to new networks, computers, operating systems, middleware, etc. In some cases, the changes can be so extensive that redesign of the system may be necessary.

Conversely, through proper use of standards and of isolation layers that hide implementation details, it is possible to design components so that application source code can be ported across a wide variety of underlying computing technologies. Achieving this objective requires exercising sufficient design and implementation discipline to forego use of vendor-unique features. This is sometimes difficult since vendor-unique features may confer a modest or even substantial advantage in performance or in initial cost; however, the long term cost of repeated use of non-portable source code to a succession of technology bases often eventually far exceeds any initial cost savings or performance gains.

Portability is not an easy goal to implement. Different types of portability are available; each resulting in a unique set of measures and metrics that is used to determine the portability effectiveness in the design. Portability types are as follows:

- *Product Line Portability*: Vendor upgrades its process such as clock speed, additional memory, etc but is nothing more than an extension of the vendor's product line in the same family (No changes to operating systems, compilers, linkers, debuggers, etc)
- *Development Environment Portability*: Vendor upgrades compilers, linkers, OS, patches, tools, processes, etc.
- *Performance Portability*: The performance in one computing and networking environment is the same (within thresholds which needs to be defined)
- *Operating System Portability*: Portability of applications between operating systems
- *Platform Portability*: Portability of applications between different platforms. This is both between a PowerPC from Vendor X to a PowerPC from Vendor Y and PowerPC from Vendor X or Y to Intel Pentium from Y or X. However, this portability is still in the basic platform style such as single board computers (SBC), i.e., SBC-to-SBC but not SBC-to-server
- *Computing Environment Portability*: Portability of applications between different computing and networking environment. This involves porting the application from one type of environment to another style, e.g., SBC-to-server

- *Language Portability*: Portability of the application from one computer language to another language.
- *Architecture Portability*: Port or convert the architecture style to a new architecture style.
- *Intellectual Portability*: Portability of knowledge and training from one environment to another
- *Complete Portability*: Combination of the different portability types or all of the portability types

A complete discussion of portability should address most, if not all, of these attributes. In many cases, metrics will involve thresholds for labor efforts, cost, schedule, and code changes to support the desired portability. Establishment of specific metrics for each of these is beyond the scope of this document and should be addressed early in OA design.

#### **3.2.2.2.2 Guidance**

In accord with the objective of producing application components that can be ported quickly and affordably from one computing environment to another, applications should be built to be portable across a variety of standards-based platforms, networks, operating systems and other technology components.

1. Application source code should require a minimum of compilation modifications due to changes in network configuration and technology, computers, operating systems, middleware, and other technology base components of the OA computing system.
2. In support of the above guidance, the changed computing environment components should adhere to the interface standards employed by the OA application software (see Section 3.2.3 for information on the OA computing environment; see Section 3.2.7 for guidance on standards selection).
3. There should be a minimum of dependence of application components on vendor unique features. Where needed, tailoring to specific machine/OS/vendor differences should be accomplished by isolation of changes in an adaptation layer. When possible, such mechanisms should utilize commonly available functions that can be emulated by other means (such as macros/stubs).
4. Portability may on occasion need to be sacrificed (e.g., via the use of vendor unique features) to meet specific component performance requirements. The Navy should approve reductions in the level of required portability. The Navy should also approve any use of vendor unique features.

### **3.2.2.3 Location transparency**

This section provides rationale and guidance for designing OA components to ensure location transparent allocation and operation.

#### **3.2.2.3.1 Description**

A fundamental and recurring problem in developing and evolving large systems is the need for frequent change and growth. This means that task-to-resource bindings established during initial system design are likely to change over time – sometimes quite rapidly and extensively – as new requirements emerge and are incorporated. System designs that fix the task-to-resource bindings at the design stage are inherently difficult, expensive and time-consuming to change.

Mechanisms exist to defer binding of application programs to system resources until system startup and, increasingly, at run-time. Use of these deferred binding mechanisms, whenever tactical performance requirements permit, confers considerable flexibility in the design and maintenance of large, complex systems. Deferred binding permits incorporation of new computing resources to meet expanded requirements while preserving flexibility to modify system resource allocation, as well as to provide more flexible failure recovery schemes. These mechanisms sometimes add overhead, a fact that has weighed against their use in the past. However, the low cost and speed of modern processors and networks renders this overhead affordable in all but the most stringent performance cases.

In accord with the objective of producing designs that permit or enable deferred bindings, applications should be built to be portable across a number of platforms, networks and operating systems. They should take advantage of modern middleware and other technologies, including collocation facilities, so that they are location transparent (i.e., so that they communicate as logical entities rather than fixed-mapped entities located at pre-defined physical addresses). Included in the capabilities of such middleware products are name services that provide the translation from the logical addressing scheme of the middleware-based applications to the underlying physical addresses characteristic of the equipment base.

Location transparency has certain implications on system test and certification. In particular, it creates a situation rarely faced in the tactical community before, namely the fact that a particular functional capability may be achieved by multiple (identical) physical allocations. (In the more general case, even this physical isomorphism may be relaxed.) While these configurations are logically identical and physically comparable, they are not physically identical. The issue of acceptance testing and certification should be addressed in this context. Certification is discussed in Section 3.2.6.

#### **3.2.2.3.2 Guidance**

To achieve the objective of producing designs that permit or enable deferred bindings, the following guidelines are provided for component design:

1. Application components should be designed to be location transparent within OA network (i.e., there should be no application-level binding between components and

physical addresses). This guidance should, at a minimum, apply to groups of compatible resource units (i.e., pools of resource that are effectively interchangeable from the perspective of component resource requirements).

2. Except where unique performance-driven configuration requirements exist, addressing of components within the OA implementation should be by means of logical names rather than physical locations. High-level middleware protocols (i.e., above the level of the Transmission Control Protocol/Internet Protocol (TCP/IP) and the Universal Datagram Protocol (UDP/IP)) that provide logical naming services should be used for purposes of inter-component communications in all cases where performance requirements permit their use. These middleware protocols should provide a naming service.
3. Object collocation facilities should be used where necessary as a means of increasing efficiency and performance.
4. Currently, complete separation of components from physical addresses is infeasible in most I/O device driver designs. To minimize the impact, one option is that the device driver should provide an interface to the system via a dedicated application component or components that present a named interface "port" or mechanism (e.g., Input/Output Processor [IOP] Gateway or storage-related nodes). Other implementations are possible and may be employed so long as location transparency is not compromised, subject to item 2. above.

#### **3.2.2.4 Client-Server**

This section provides rationale and guidance for designing OA components according to the widely used client-server design pattern.

##### **3.2.2.4.1 Description**

The client-server design pattern has gained widespread acceptance in commercial computing applications. Indeed, it has become almost ubiquitous because the model provides a very useful way of exporting services to a variety of users in a standardized manner. Servers, in particular, provide an excellent means of promoting reuse and standardization at the application level.

A server is a system component that supplies and controls access to a computational or data resource via a predefined interface. The resource controlled by a server may be a physical device, such as a display or external sensor; a persistent data store, such as a database; a state data repository; an algorithm or a combination of these.

The interface provided by the server may include specification or guarantee of QoS parameters, such as client request priority or maximum response latency. A server exhibits the following characteristics:

- A server and/or clients may consist of multiple cooperative or replicated processes.
- Servers do not require any architectural or functional knowledge of the client application(s) that use it.
- Clients require no functional knowledge about the server other than the server's API; nor

do clients need to know the specific server replica or process with which they are communicating.

- A server does not initiate interaction independently with its clients unless clients have registered for specific “push” services (i.e., server-initiated updates).

A client is any component that accesses the server via the defined API. Consistent with previous discussions, clients and servers may be physically collocated where use of such mechanisms is appropriate for performance. A server has no a priori knowledge of its clients; thus the client should initiate any interaction between client and server. Similarly, clients do not need to know the internal workings of the server. Furthermore, since a server is not dependent upon the function or implementation of a client, it may provide services to multiple heterogeneous clients.

Because clients and servers may be mission critical components, and because they may have significant processing loads and/or throughput requirements, both clients and servers may have replication requirements in support of fault tolerance and/or scalability. Replication of any given client, for either reason, should require no modification of the predefined client-server interface. However, the interface can be designed to provide server support for client replication without negatively affecting the integrity of the client-server relationship. Examples of interface design choices that would support client replication, while still maintaining an appropriate decoupling between client and server include:

- A registration mechanism in which the client specifies a publication group for the server to use in its distribution of data to the client. All replicas of the client can then benefit from the server’s data distribution without each replica having to register separately.
- An interface by which a new replica can request current state data from the server. By anticipating the requirements of a replicated client, an appropriately designed server interface can work equally well for replicated and non-replicated clients.
- Interface support that allows a client to register for only a selected portion of data (e.g. only certain tracks) for the purpose of load sharing with its peers.

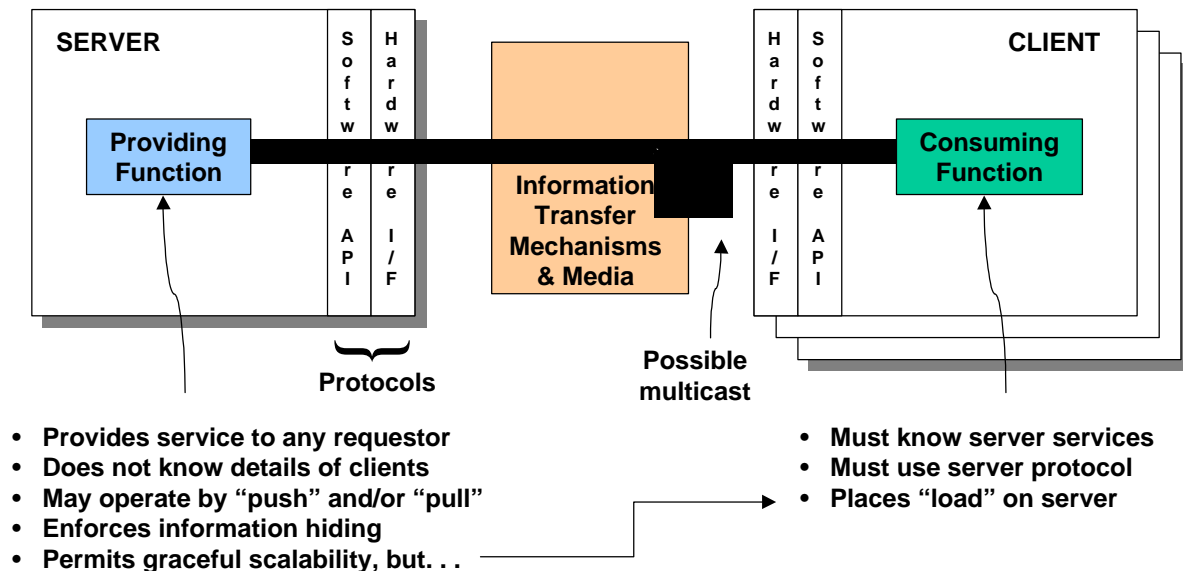
The roles of client and server are not mutually exclusive. These designations refer to the relationship of two processes with regard to a set of interactions, e.g. access to a database. It is possible, and even likely, that a given process can be a client with respect to one set of interactions, and a server with respect to another. For example, a server that distributes kinematic sensor data may also be a client of a server that distributes non-kinematic sensor data – such as an identification server. In addition, both servers may be clients of a data-recording server.

Use of client-server architecture concepts allows for the development of both subsystem-wide and system-wide services. Typical services include but are not limited to data distribution; access to physical resources, such as sensors, actuators, or external devices; or state data protection and mediation. They may provide services required by the entire system, such as time synchronization or instrumentation; or that may be required by only a subset of the applications. Regardless, the availability of client-independent services greatly enhances the ease with which new elements and applications can integrate with the existing system.



A schematic of the interactions between clients and servers is shown in Figure 5. Three common models of client-server interaction are defined:

- Synchronous: request-reply
- Asynchronous: one-way
- Asynchronous response: client registration



**Figure 5. Client-Server Model**

The first model, the synchronous model, is also known as the pull or request-reply model. This model requires that the server take a specific action in response to the client’s request. A client will issue a message to a server requesting that a service be performed. For example, the client might request the display of data at an operator console or the initiation of some algorithmic processing. When the server completes action on the request, it will respond to the client with the results of the action, i.e. with an acknowledgment of completion or the results of the computation.

In some cases, a client may make a request to the server and no response is required. This type of interaction is the second common model, sometimes called the asynchronous one-way model, when the results of an asynchronous request to a server are implicit (e.g., a change of state data in the server).

The third common model of client-server interaction is the client registration or asynchronous response model, often referred to as the push model. This model requires that clients register for services that are to be delivered at a later point in time. For example, a client may register with a track data server for the update of track data at some client-specified frequency. Alternatively, a client may request notification upon a change in state of some resource. The server retains information about the specific requests of the registered clients as well as sufficient logical addressing information to allow the server to monitor the status of the clients and provide the information as appropriate.

Peer-to-peer relationships are typically used between process replicas or between related clients of a common server. Peers communicate to notify one another of state changing events or to coordinate a set of activities that culminate in a change in state data.

Servers control development and maintenance cost and complexity by providing software that serves many clients instead of providing separate and unique software functions for different consumers. Thus, the client-server model supports the concept of multipurpose software. If unique customization were permitted, the client-server design pattern would not be helpful.

However, a critical system performance impact exists that should be accounted for during server design. By their very nature, clients place loads on the servers they access. Thus, servers should be designed carefully to ensure required system performance and efficiency. For this reason, design of client-server architectures, particularly for real-time applications, should be approached carefully. Furthermore, care should be taken that “utilization creep” does not occur; i.e., a growth in the number of calls on a server pushes the server towards the upper limits of its planned load capacity without sounding any alarms. In order to mitigate the impact of this client-imposed load, it is appropriate to consider a scalable design pattern for servers (see section 3.2.2.9).

#### **3.2.2.4.2 Guidance**

The client-server model has been shown to be beneficial in implementing flexible and scalable systems. As such, it should incorporate the following guidance as appropriate:

1. To the maximum extent possible, servers should be designed such that they are independent of the implementation of their clients. Servers should be capable of providing services to multiple heterogeneous clients.
2. number of common models of client-server interaction are available to component designers. One possible model is the synchronous or request/reply model. Another model is the one-way or asynchronous model. A third model is the client registration, or asynchronous push model. Other models are possible as well.
3. Clients and servers may have replication requirements in support of fault tolerance and/or scalability. Replication of any given client, for either reason, should require no modification of the predefined client-server interface. However, the client-server interface can be designed to provide server support for client replication without negatively affecting the integrity of the client-server relationship.
4. For clients and servers where replication is involved, the client-server API should be designed such that the clients are not directly dependent upon the replication mechanisms of the server. Additionally, the API should support, or at least not interfere with potential replication of both the servers and the clients.
5. As with all component interfaces (see Section 3.2.2.1), the interface provided to the client by the server should be well defined in an appropriate interface specification document or via an Interface Design Language (IDL). The API for this interface should be defined in an acceptable language or meta-language (e.g., Ada package specification, C or C++ header file or CORBA IDL). It should provide access to the server via a procedural or object class interface. When the API and server are separately maintained, documentation should be required of both the API and the public interface between the

API and the server.

6. Servers should detect erroneous invocations and invocations for which sequence dependencies exist, and they should export exception conditions to invoking clients via the component's interface. Such erroneous invocations should not cause the component to operate incorrectly. Interface specifications should indicate the meaning of exceptions and the conditions under which they arise (e.g., the specification should define what happens if **open()** is called twice without an intervening **close()** call).
7. Components should adopt internal design safeguards (e.g., finite state machine structure), to eliminate the possibility that erroneous inputs and/or call sequences can result in erroneous operation, hung conditions and crashes of the component.
8. Definition of server interfaces should be performed by the organization responsible for server development; however, potential client development organizations should be involved in the review process prior to server interface finalization. Where feasible, server functional enhancements may be incorporated via added API functions rather than changing the syntax or semantics of an existing API function. Obsolescence and eventual discarding of API functions should be addressed as a life-cycle issue.
9. Where appropriate, the server interface specification should include both explicit and implicit QoS parameters (see discussion in Section 3.2.2.10.1.2). An example of an explicit QoS parameter would be a mechanism for clients to specify response latency with their request. An example of an implicit QoS parameter would be a maximum response time associated with a particular type of request – regardless of the client initiating the action.

### **3.2.2.5 Data Distribution**

This section provides rationale and guidance for ensuring the efficient flow, or distribution, of high-volume continuously refreshed data among OA components. Examples include distribution of track updates and navigation data to OA components. Distribution management techniques include use of a track data server to provide track distribution to registered track consumer client components; use of periodically refreshed push distribution (forward caching) of continuously refreshed data; and use of asynchronous processes with timed/fixed packet I/O buffering.

#### **3.2.2.5.1 Description**

The nature of data flow within a large, complex system often dictates that certain components are distributors of data that is continuously refreshed, frequently in accord with a periodically scheduled thread that guarantees that no data are ever more stale than a defined value. Such components are excellent candidates for use of the client-server design pattern. In order that they perform as effectively and efficiently as possible, certain design characteristics should be implemented.

Many tightly timed real-time systems resort to a highly rigid synchronous periodic design in order to achieve efficiency. This is a good design for pure hard real-time systems that cannot meet timing requirements any other way. However, it is also an inflexible one, particularly when applied in a distributed processing environment. If timing and performance considerations

permit, a far more flexible approach is available; i.e., one based on asynchronous loosely coupled processes.

In such cases, workload flow variations are generally “cushioned” by use of message queues or buffers to smooth out backlogs and variations in arrival rates. This approach can be used in conjunction with the synchronous, hard real-time approach, thus fostering co-existence of hard and soft real-time processing within the same overall system.

Managing the flow of frequently updated information in such a real-time system is a critical design characteristic, particularly if the flow is high volume, as it so often is where sensor data distribution is part of the overall system design. This problem is compounded if the updates are irregular, as they are in the case of some sensor systems. Such variations are attributable to the fact that data flow is a function of what is in the environment in which the sensor operates.

In real-time systems, requirements generally tend to focus on latency considerations, despite the fact that -- in view of the high volumes associated with these data streams -- efficient use of available bandwidth is also an important consideration. The two simplest strategies for managing this data flow are (1) to send every update, which minimizes latency, or (2) to fill a fixed-size buffer prior to transmission, which maximizes channel throughput.

However, neither strategy provides a complete solution in many real-time system applications. In particular, the delivery time for the fixed-buffer method is potentially unbounded, an unacceptable situation for real-time performance. A good compromise is to employ both techniques simultaneously.

In the combat system, track distribution is a good example of the data distribution problem, due to many consumers of track data. In previous system designs, processed track information (i.e., correlated and filtered data) was provided to consumers via custom designed interfaces, with each interface being uniquely tailored to the requirements of the receiving element. This provides maximum functional capability, but it is also the most costly approach, and the one least susceptible to multi-purpose use in the client-server sense.

Providing push-based server distribution of the entire track file, appropriately time-tagged for local extrapolation (see Section 3.2.3.13 for guidance on system-wide time synchronization), to all track consumers ensures that track information is always available to consumers when they need it. In effect, data is “forward cached” with track consumers against the certain eventuality that it will be needed. The volume of data and the frequency with which it should be distributed is well within the capacity of modern switch-based networks.

Benefits of using a push-distribution mechanism for track data include reduced processing in the track data server, reducing the possibility of widespread contention and performance bottlenecks. There are also the benefits of increased ease of integration of new functionality; ability to leverage common code to perform local caching and retrieval of track data, resulting in a reduction of application complexity; minimal performance and bandwidth impact with additional track data consumers if multicast is leveraged.

### 3.2.2.5.2 Guidance

The design guidance for OA components that perform distribution of continuously refreshed data is as follows:

1. The client-server asynchronous response model (see Section 3.2.2.4) should be used for distribution of continuously refreshed data. While this would seem to drive up network bandwidth utilization, techniques such as multicast exist to ameliorate this possibility.
2. Servers should have the capability to accept registration requests from client components.
3. A common API and library of services should be provided by data servers to data consuming clients for receipt of events and retrieval of locally cached data.
4. Data distribution servers should provide a defined class of service with respect not only to functionality but also to performance. Thus data distribution periodicity, senescence and maximum load should be defined as a component requirement for the server. The periodicity and other distribution characteristics of the server's transmissions should meet consumer requirements. Distribution servers should be designed to provide a limited but distinct range of distribution QoS options in order to service varying data consumer timeliness requirements.
5. Where throughput and latency requirements require it, continuously refreshed data flows among tactical components may be implemented using timed buffer transmissions. In this design, the buffer size is often fixed, but its transmission is tied to a periodic cycle that satisfies timeliness requirements associated with the data flow. If the buffer fills up before the period expires, it is transmitted as soon as it is full. In this way, efficient throughput and channel use is maintained. On the other hand, if the buffer is not full when the period expires, then the partial buffer is transmitted. This ensures that a defined worst-case latency is not exceeded.
6. To ensure efficient server operation, and when possible given performance requirements, underlying communication of data from data distribution servers to their clients should be performed by middleware protocols that provide IP multicast functionality. This ensures efficient server operation.
7. Where applicable, data distribution servers should be capable of guaranteeing first-in, first-out order of data flow to each subscribed client. However, total ordering of server data flow across clients is not necessarily required. That is, unless specific applications have a total ordering requirement, completion of a particular update to all clients is not necessary before some clients begin receiving other updates.
8. Delays in consumption of distributed data on the part of a client should not result in a system-wide performance-degrading backup of data distributed from the server to that client or any other client; nor should the server's operation be impacted. Where required, clients should be designed such that, during overload conditions, data may be discarded. This latter requirement ensures that the ordering requirement stated above does not force clients to consume stale data.
9. Mission critical data distribution servers should be fault tolerant. A number of useable fault tolerance strategies are defined as architecture patterns in Section 3.2.2.8.
10. Where substantial growth is anticipated in the load a data distribution server is likely to experience over its life-cycle, the server should be designed using a scalable architecture pattern (see Section 3.2.2.9).

### **3.2.2.6 State Data Coherency**

This section provides rationale and guidance for ensuring that OA components that manage composite state data are designed to provide correct operation under all possible state transition situations. Included are complex state variables resulting from the run-time interaction of many components. Maintaining state data coherency is a part of the overall requirement to preserve data integrity in mission critical and safety critical systems.

#### **3.2.2.6.1 Description**

Certain applications are inherently state intensive. For such applications, the total set of state information constitutes a composite state vector comprised of state information contributed by many other components. Maintenance of this composite state information is often vitally important for correct system operation. Furthermore, maintenance of and interchange of state information among replicated components for the purpose of failure recovery is a major consideration in designing state-intensive fault-tolerant components. For this reason, special consideration should be given to the design of components to ensure state data coherency.

Mechanisms, by which state data coherency may be maintained include atomic multicast communication transactions with one or more components, guaranteed delivery of message traffic in a specific order and synchronized exchange of state data among participating components. Atomicity of multicast communications means that for each multicast, either all recipients receive the message or none receive it – and that the sender knows who received what. In the case of ordered messaging, various ordering properties are frequently defined (see Section 3.2.3.7.1.3).

Use of priorities can, at times, conflict with maintaining state consistency. Priorities should not be used to achieve message ordering, since the scheduling mechanisms between platforms vary and can result in different effects for applications running on different operating systems. Also, priorities can disrupt other ordering mechanisms, such as ordered messaging, potentially causing the order in which messages are handled to be unpredictable.

#### **3.2.2.6.2 Guidance**

The design guidance for OA components that maintain composite state data is as follows:

1. Where application performance requirements permit, components that manage composite state data should utilize middleware and reusable services to enforce communications atomicity, message ordering and state data exchange. Impact on application source code should be minimized and limited to those steps for which appropriate middleware and service library functions are not available.
2. Components with composite state requirements should not utilize priorities to achieve ordering between replicas. To the degree possible, components that maintain composite state should not depend on use of priorities for correct operation. Where essential to meet performance requirements, care should be taken when using priorities in such components so that affects on portability and interoperability are minimized.
3. Components whose primary task is state data management should ensure by component

design that computations associated with state data coherency maintenance do not adversely impact performance and timing characteristics of other system components. Maintenance of composite state arising from interactions with other components places a strong obligation on the designer to ensure responsiveness of the state data-maintaining component. Further investigation is required to define specific mechanisms.

4. Unlike data distribution servers (see Section 3.2.2.5), which have a requirement for non-coupled operation with their clients, state data intensive components should operate in synchrony with other components to which it interfaces to maintain its composite state.
5. State data maintenance should be compatible with the component's fault tolerance design pattern (see Section 3.2.2.8).

### **3.2.2.7 Computational Flow**

This section provides rationale and guidance for ensuring those OA components that perform periodic or on-demand computations are designed to provide efficient and effective operation.

#### **3.2.2.7.1 Description**

Many of the computations in large, complex systems are part of an overall data flow process. In this model, information arrives at a component, is processed, and is then forwarded to another component for additional processing. Such components are sometimes referred to as pipes or filters. The concatenation of several such components, coupled by interfaces and by component-spanning timelines requirements, constitutes a data flow path through the system.

Data flow paths can include both periodic processing and event-based processing. They may also be quasi-continuous or intermittent.

For periodic data flow paths, the nature of the processing is such that large quantities of data are processed on some pre-defined periodic frequency. In some instances, outputs may occur only infrequently (e.g., when certain conditions are detected in the incoming data). In other situations, periodic outputs of processed data may occur. Periodic data flow processes should also execute in a coordinated fashion with external event reaction time paths. Since event-based processing should execute with very low latency, periodic data flow processes may need to be pre-empted, requiring that special care be taken in their design.

Event-based processing is defined as processing that takes place in response to some pre-defined external event or condition. Sometimes the resultant processing is limited to either a discrete response to the event or to updating internal state information based upon the event. In other cases, a quasi-continuous sequence of processing may be initiated. In the quasi-continuous cases, a computational data flow may come into existence for a period, after which it returns to a quiescent state.

#### **3.2.2.7.2 Guidance**

The design guidance for OA components that perform computations as a part of an overall system data flow is as follows:

1. Components that are part of a computational data flow path should, where necessary, utilize priorities to ensure that reactive threads of operation requiring minimum latency response are able to pre-empt routine threads to gain control of the CPU.
2. Where routine computational threads (whether periodic scheduled or externally stimulated) are coupled to physics-based system requirements, thread priorities and careful system design should be utilized to ensure that processing requirements and latencies are achieved.
3. When quasi-continuous processing requirements may occur, the system designer should ensure that adequate reserve computational capacity is designed in to ensure proper performance of both continuous and quasi-continuous operations. Use of a dynamic RM capability to ensure that component computing capacity requirements are met is discussed in Section 3.2.3.8.
4. To the maximum extent possible, computational data flow components should be constructed to minimize retained state information. Unless otherwise infeasible, state data maintenance for such components should be limited to state information that is needed to sustain fault tolerance and failure recovery operations.
5. For components that should maintain state data, every effort should be exerted to design such components so that state data updates consist of a simple update operation involving only the previous state operated on by new data. Interactions with multiple other components in the maintenance of state data should be avoided to the maximum extent possible. When these conditions cannot be met, then state data coherency preserving patterns are likely to be more appropriate (see Section 3.2.2.6). Further investigation is required to define specific mechanisms.
6. Computational data flow components for which computational demand is likely to exceed the capacity of a single CPU - either in bulk computational capacity or in required minimal computational senescence - should be designed to exercise computational load management among a number of components (see Section 3.2.2.9). These components may be scalable peer-clients, pipelined components, partitioned functional capabilities, etc.

### **3.2.2.8 Fault Tolerance**

This section provides rationale and requirements for ensuring that OA components requiring continuous availability are designed to provide a fault tolerant capability in accord with a defined set of fault tolerance models (e.g., active, passive and selected hybrids).

#### **3.2.2.8.1 Description**

Failure of a mission critical application process means that the system can no longer carry out an assigned function. When rapid recovery is required, replication of mission critical application processes can increase the tolerance of the system to such faults. This is known as replication for fault tolerance. When long recovery time specifications exist, replication for fault tolerance may not be necessary. In such cases, the system RM function may simply restart the failed application component to regain the functionality that was lost as a result of the failure.



In addition to replication for fault tolerance, components may also be replicated for load sharing (see Section 3.2.2.9). In general, load-sharing components also contain provisions for readjustment of loads and continuation of operations following faults.

Frequently, mission critical processes retain memory-resident critical-system state data. As such, in the event that a replica fails, maintaining the integrity of that state data is a significant consideration in the design of an application that is replicated for fault tolerance.

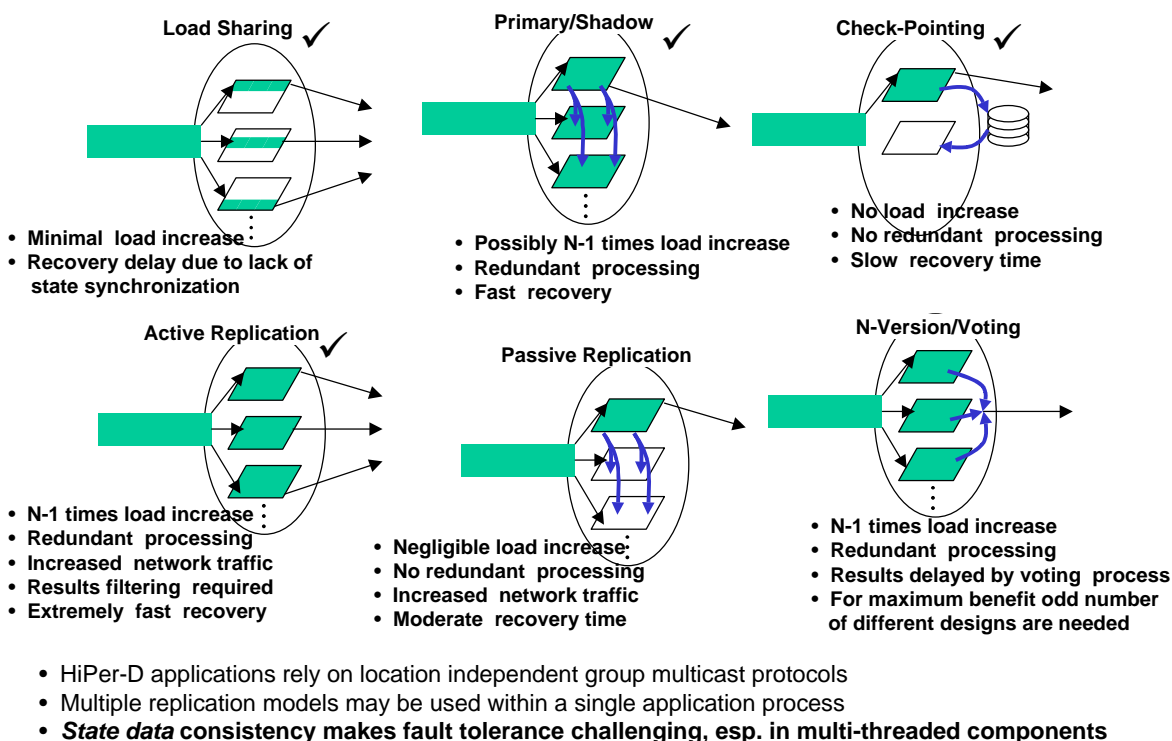
Middleware plays a critical role in the implementation of any fault-tolerant component, especially group communication middleware. Group communication middleware, including communications endpoint considerations, is discussed in Section 3.2.3.7.1.3.

Several common models used to implement fault-tolerant replication solutions. Six models are illustrated conceptually in Figure 6, of which four are described in more detail in the following subsections - Active Replication, Passive Replication, Primary/Shadow Replication, and N-Version Replication. The load-sharing peer-client model is discussed in Section 3.2.2.9.1.2., and Checkpointing will not be discussed as it is a widely used and well-understood technique applied primarily to non-real-time systems. The items that are checked in Figure 6 – Load Sharing, Primary/Shadow, Check-Pointing, and Active Replication- represent fault tolerance models examined in the HiPer-D risk reduction program.

**3.2.2.8.1.1 Active Replication.** When an application process has deterministic inputs/outputs and deterministic behavior based on those inputs/outputs, active replication is a valid model for achieving fault tolerance. Active replication involves executing multiple identical copies of the same replicable process. Each replica receives all inputs, performs all processing, and provides responses. Receivers of responses from the replicas are responsible for filtering out duplicate data contained in the responses. The active replication model relies on deterministic input and behaviors for the replicas to remain in a consistent state. Generally, state synchronization between replicas is not performed except during replica initiation.

Two significant limitations are associated with this model: first, very few applications are completely deterministic. Sources of non-determinism may include very complex processing associated with reliable and atomic multicast, multi-source data inputs, memory management, thread scheduling, distributed time synchronization, time-based computations, and error handling. Group communications tools (discussed in detail in Section 3.2.3.7.1.3) can be very useful insofar that they can provide atomic, totally ordered multicast and ordered group membership change notifications. These attributes can reduce or eliminate potential non-determinism in the input/output streams to the replicas; however, such tools do not control other sources of non-determinism.

The second limitation with active replication is its inability to handle Byzantine failures. This class of failures is extremely difficult to manage with any model other than N-version



**Figure 6. Fault Tolerance Models**

replication, discussed in Section 3.2.2.8.1.4.

Single-threaded applications that perform non-time dependent calculations and applications that perform database retrievals from non-writable sources may be able to utilize an active replication model; however, its applicability is limited in the complex multi-threaded real-time systems domain.

**3.2.2.8.1.2 Passive Replication.** Passive replication is an appropriate replication strategy when the behavior of a process is non-deterministic based on its inputs and/or when processing resources are constrained. As in active replication, all passive replicas receive all inputs. Only the primary replica actually processes inputs and provides responses. The other replicas referred to as the secondary or passive replicas, queue the inputs until checkpoint data is received from the primary. The replicas use the checkpoint data from the primary to update state data as required, and the queued inputs associated with that checkpoint data are deleted. When a primary replica failure occurs, one of the passive replicas assumes primary responsibility and begins processing the inputs that were queued since the last checkpoint was received.

This method of replication requires associated applications be capable of handling duplicate data from the replicas. Recovery time may be long, if significant processing is required for the queued inputs. To some extent this is controllable. Recovery time is determined by the

relationship between queue growth and checkpoint interval. However, this model tolerates non-deterministic responses to input better than other replication methods.

Clearly, successful implementation of this replication model requires that associations be established between the inputs processed by the primary replica and the checkpoint data. By ensuring totally ordered message inputs to all replicas (both passive and active), group communications tools can play a significant role in establishing these associations.

**3.2.2.8.1.3 Primary/Shadow Replication.** A significant number of mission-critical applications do not meet the determinism criteria to use an active replication model, yet they have real-time requirements that cannot be fulfilled through passive replication. For these applications, a primary/shadow replication model may be more appropriate.

Although replicas processing the same inputs and outputs may arrive at different solutions, the primary/shadow model assumes that both solutions are potentially correct (in the absence of various subtle hardware and software errors). Multiple copies of a process are executed simultaneously, with one designated the primary replica and the remainder designated shadow replicas. All replicas receive and process all inputs with the primary replica providing control over the order in which the inputs are processed. When the primary arrives at either an intermediate or final solution, it provides that solution to the shadow replicas. The shadow replicas update their state with the primary replica's solution, and do not issue independent responses, thus maintaining consistency with the primary replica. If the primary replica should fail, one of the shadow replicas assumes the role of the primary at the point in the input/output flow at which the failure occurred.

This replication model is based upon the following assumptions: failures are ordered with respect to message flow, messages will be issued via atomic multicast, and messages will be processed in first-in-first-out (FIFO) order. Since the primary replica controls the order in which the inputs are processed and does not require identical responses from the replicas, multi-threaded applications can be handled, and totally ordered messages are not required.

A detailed understanding of potential sources of non-determinism in the application is required to use this replication model. Since these sources vary with the specific application, primary/shadow replication cannot be implemented without awareness of the synchronization of the application state between the replicas.

**3.2.2.8.1.4 N-Version Replication.** For some applications, correctness and/or uninterrupted (seamless) availability are critical. Safety critical systems have unique correctness requirements. One method of providing greater assurance of algorithmic correctness is obtained through utilizing multiple implementations of critical algorithms. In N-version replication, N implementations of a process are executed, all receiving inputs and generating solutions. Through some arbitration, such as voting, a response is chosen, and potentially misbehaving processes are identified and excluded.

This method of replication shares some similarities with active replication. Care should be taken to exclude or accommodate sources of non-determinism, such as time dependencies. Some debate exists as to whether N-version replication provides significant benefit with respect

to software/algorithmic correctness, since the source of many software and algorithm errors lies in the specification of the algorithms. Additionally, many software errors also result from incorrect handling of boundary and error conditions; thus they are more likely to occur in multiple implementations. Due to its high cost, use of N-version replication is rare.

**3.2.2.8.1.5 Client-Server Interactions.** Examples of how a client-server interface can support client replication include:

- The client-server interface may be designed to allow the clients to readily subdivide data distribution among multiple peers, simplifying the ability to implement scalable peer-clients.
- The client-server interface may allow the clients to specify a multicast address for data distribution or for request results, supporting replication techniques that provide full availability by instantiating duplicate running copies.
- The client-server interface may allow a state transfer to occur when a new member subscribes, simplifying the initiation of new client replicates.

#### **3.2.2.8.2 Guidance**

The system should be resistant to faults in processors, network components, system software, and application computer programs. Component replication should be the primary mechanism for achieving fault tolerance. The guidance is as follows:

1. All OA time and mission critical components should be replicated for fault tolerance. In this context, time critical is defined as components for which reloading from a backup media is not adequately responsive. Components should be designed and implemented in accord with the appropriate replication strategy and fault tolerance design pattern.
2. OA components that are not time critical may be designed to be restartable under the control of the system dynamic RM capability. For such implementations, Checkpointing may be a viable state data restoration strategy. Other strategies for fault tolerance, including those applicable to time critical components, may also be used.
3. Under normal conditions, each load sharing replicated component will actively share the processing load for the particular function. A minimum of two instantiations (or replicates) of any process on a mission critical path should be executing to support fault tolerance.
4. Metrics for assessing fault tolerance should be defined and performance thresholds should be established. The metrics will be used by the RM capability (see Section 3.2.3.8) to manage resources. System degradation should be measured in terms of task completion, timeliness and data integrity at a minimum. Further investigation is required to define specific mechanisms.
5. Detection of a fault in a processor, network component, or software component that results in loss of the component's capability (e.g. a fault that is not masked by network redundancy) should result in automatic shifting of the component's processing load to a replicated component. This may be a shifting of load to already active instantiations or may require that additional instantiations be replicated and activated to facilitate load

redistribution. Detection and recovery times should be equal to or better than existing NWS requirements.

6. Each tactical computer program component should be independent of other components with regard to fault tolerance to the degree practicable. A failure in one component and the subsequent recovery actions associated with restoration of the capability should not require that other tactical components be replaced in unison except in limited situations when dependencies have been determined to be unavoidable. Each non-failed component should respond properly as the failed component recovers. Load sharing provisions are discussed in Section 3.2.2.9
7. Multiple simultaneous component failures (e.g. loss of a cabinet containing multiple processors) should be treated as a series of independent failures. The failure management capability should queue all pending recovery actions and order them to ensure that, based upon recovery dependencies, multiple recoveries of the same component do not occur.
8. The fault-tolerant implementation should be designed to execute in a unified fashion across all OACE components. That is, for application components utilizing multiple OACE technologies, the varying characteristics of OACE shall be considered in component design. The use of fault tolerance features from multiple OACE technologies should not adversely impact application component fault tolerance performance with respect to fault tolerance metrics.
9. Due to its high cost, use of N-version replication should only be used when absolutely necessary.

### **3.2.2.9 Scalability**

This section provides rationale and guidance for ensuring that OA components with high volume and large dynamic ranges of computational demand are designed for scalability.

#### **3.2.2.9.1 Description**

The term “scalability” has multiple meanings. In summary, these meanings of scalability are as follows:


- Resource scalability
- Functional scalability
- Load scalability of load sharing

The first form of scalability denotes the ability of the system’s design and implementation to accommodate incorporation of additional computing resources. Hence, networks promote scalability by making it easy to add additional computing resources.

The second type of scalability is the ability to add new functional capabilities without disrupting existing functional capabilities. The client-server model promotes functional scalability by allowing new client functions to access existing servers without disrupting other clients already using the server.

Finally, application components may themselves be made scalable to varying processing and input/output loads via architectural design patterns and load-sharing mechanisms. It is this last form of scalability that is discussed in this section. Two methods for implementing scalability are discussed, functional partitioning and replication.

**3.2.2.9.1.1 Load Scalability by Functional Partitioning.** Load scalability may be implemented by sufficiently partitioning a component's processing algorithms so that the resulting executable processes can be spread across a number of processors and nodes. The partitions should be selected carefully so that the system can withstand, with a reasonable margin of safety, the maximum anticipated load. Common partitioning patterns include pipelining and paralleling. These patterns are shown in Figure 7. In the upper part of the figure, the boxes (p1, p2, etc.) represent pipelined components. Each process processes its inputs and passes outputs on to the next component in the pipeline. In the lower part of the figure, each of the parallel functions (f1, f2, f3) processes data of a particular type; different from data processed other functions.



**Figure 7. Partitioning Scalability Patterns**

Pipelining is used when multiple separable computations should be performed serially on a set of continuous updated data. It is most effective when the computational and input/output load has a clear upper boundary. Typical applications include sensor data processing. In pipelining, computations are split between an appropriate number of applications. Ideally, each application requires an approximately equal time to process each set of data. This maximizes the utilization of the processors. If one computational application requires significantly longer than the other applications in the pipeline, it may introduce a performance bottleneck that interferes with the throughput of data. As each application finishes its computations on a given set of data, it passes its results to the next application in the pipeline and retrieves the data from the previous

computation. The result is a relatively continuous, ordered flow of data and processing that can be spread across multiple processors.

Paralleling is applied for components that have computations that lend themselves to simultaneous execution. One application distributes the necessary data to multiple computational clients. When all computations are complete, a server that represents the endpoint of the processing may synchronize and synthesize the resulting data as required. As in pipelining, paralleling behaves optimally if all computations require roughly equal time, particularly if synchronization of the computations is required at the endpoint. An example of paralleling includes doctrine processing, where each track can be qualified for each of the different types of doctrine and each of the active doctrine statements simultaneously.


Combinations of architectural patterns may also be used to enhance scalability. For example, a component may use paralleling, while one of the parallel paths comprising that component uses pipelining. Additionally, each component in the partitioned component may use other architectural patterns that support other key design characteristics. For example, a component in a pipeline may use an active replication pattern for fault tolerance.

Scalability through partitioning provides good scalability within the limits of the intended range of load. It does not, however, provide support for dynamically adapting to changing load. It is extremely useful where the load is known a priori to be bounded within a given range, and where there is no significant advantage in being able to reclaim resources when the load is low.

**3.2.2.9.1.2 Load Scalability by Replication.** Another method for providing scalability is through scalable process replication. With scalable replication, process replication is used to enable load sharing. Each replica of a load sharing performs a portion of the total workload. Different techniques for scalable replication may be required for clients and servers. Clients that are replicated for load sharing are called peer-clients. Servers that are replicated for load sharing are called scalable servers.

Figure 8 illustrates the use of replicated processes (R1, R2, and R3) to achieve load sharing. Each replica processes a subset of the incoming workload. Unlike the case of paralleled functions illustrated in Figure 7, where each function processed a different type of data, the R<sub>i</sub>-replicas will process different instances of the same kinds of data.

A high-volume server is a server that has a significant likelihood of causing starvation of its clients due to the number of clients, the number of requests, or the quantity and frequency of data it is distributing. When a high-volume server becomes backlogged, it has the potential to interfere with mission critical processing, either directly or indirectly through some other inter-process dependency (similar to priority inversion).



**Figure 8. Load Sharing by Replication**

The likelihood of encountering a high-volume server backlog can be reduced through the implementation of scalable servers. Scalable servers can distribute a load by a number of different methods. A method appropriate to the loading characteristics and synchronization requirements of each server should be chosen. Some methods for consideration include:

- Assigning each scalable server a different set of clients.
- Assigning each scalable server a different set of data/resources.
- Providing a coordinator that assigns each new request in an unbiased manner (e.g., round robin, or least load) to a different scalable server for processing.

Regardless of the method chosen, scalable servers should be capable of self-managing the load assigned to each replica. Typical methods for achieving this include the use of deterministic algorithms by which responsibility is determined, and assigning one replica as a coordinator for establishing/re-establishing load delegation. The method chosen often depends on the nature of the load.

Similar to high volume servers, client components can become overloaded with information that needs to be processed. Scalable replication, in the form of load sharing peer-clients, is one way to overcome this situation. Each load-sharing peer executes as a separate process. The peers divide load through the use of a load-sharing algorithm contained within the replicas. This algorithm allocates a unique portion of the processing to each peer. The algorithm should account for the number of replicas in the system at any given point in time in addition to the type and quantity of the load. This information is used to divide the processing of data in an approximately even distribution across all peers. When the number of peer replicas changes, the processing load is readjusted at each peer. Other conditions that might require load readjustment may include changes in the type of load and or significant imbalance in the load distribution. In general, minimal coordination is required among the peers beyond the adjustment of load that occurs upon peer initialization or termination/failure.

The peer-client approach can potentially increase application performance by decreasing the overall load on a single computer. Such increases in performance need not double



performance to provide benefits to the overall system. For example, the peer-client approach can be used to dampen the detrimental effects of a failure in the system by manipulating the number and placement of peers, thus distributing the impact of failures throughout the system. Peer-clients may also improve fault tolerance by increasing the number of processors that should fail before a function is lost. Finally, the ability to increase, migrate, or decrease peer components facilitates dynamic run-time load leveling as a means for maintaining optimum system-wide utilization.

Scalable replication provides the ability to adapt a system dynamically for a wide range of loads. The implementation of scalable replication can be significantly more complex than solutions that utilize computational partitioning. It can be especially complex when the component involved should maintain state data. However, scalable replication provides support for fault tolerance in addition to scalability, making it an excellent choice for mission critical components that should be fault tolerant and have significant or a varying range of load.

### **3.2.2.9.2 Guidance**

OA should choose appropriate load sharing mechanisms to achieve scalability. In particular, the following guidance applies:

1. For components whose loads are predictable and have an upper bound, architecture patterns should be leveraged that divide computational responsibility across multiple executable processes. The partitioning should be selected such that the component is capable of handling the maximum possible load.
2. Where partitioning is used to support scalability, fault tolerance patterns should be considered to provide component availability.
3. Where the loads are highly variable, scalable replication should be utilized to maximize resource utilization.
4. Scalable servers should be designed to facilitate instantiation of one or more replicas across which load are distributed. An appropriate method for distributing load across these servers should be used. Replication patterns that support state data synchronization should be considered for servers that have state data maintenance requirements.
5. To the degree practicable, scalable servers should be capable of self-managing the load distribution across the set of available replicas. They should be capable of the assignment of load and reallocation of load as the availability of replica change, without external intervention from a resource manager, except possibly for the notification of change in the number of replicas.
6. The peer-client design pattern should be employed for clients where a wide dynamic range of computational requirements makes its use beneficial.
7. In a peer-client design, the work of a client should be divided into  $n$  components known as peer-clients or replicas. For  $N$  peer-clients, each peer should concurrently perform approximately  $1/N$  of the overall load. When a peer is lost, the remaining peers should assume a larger proportion of the work in accord with a deterministic algorithm (e.g.,  $1/(N-1)$  of the load). Conversely, adding a peer to the system should reduce the load on each peer (e.g. to  $1/(N+1)$ ). Component state data synchronization may be required between peers, which appropriate design patterns for fault tolerance and state data synchronization should be evaluated.

### 3.2.2.10 Real-time Performance

This section provides rationale and guidance for ensuring that OA components are designed so that both hard real-time and soft real-time requirements are met.

#### 3.2.2.10.1 Description

NWS functions are characterized by varying degrees of timeliness requirements. These range from human-paced decision-making to soft real-time background and throughput processing to hard real-time sensor and weapons management. Each of these requirements should co-exist smoothly in an open distributed combat system..

**3.2.2.10.1.1 Real-time Processing.** The terms hard real-time and soft real-time are often used to characterize computing timeliness requirements. In distinguishing between these two terms, some authors suggest a time domain dividing line – often associated with periodic, (e.g., < 1 millisecond, < 100 microseconds, etc.). However, this approach is not sufficiently descriptive to be fully useful. For example, it is difficult to reach agreement on the precise value of such a boundary. Different application domains tend to produce different boundaries.

More importantly, many practitioners consider hard real-time to be associated with a requirement for deterministic performance characterized by guarantees of bounded response times and latencies for operations, service calls and information transfer. In this context, soft real-time systems are those for which similar tightly bounded responses are not required.

In a practical sense, such applications are most often associated with physics-based computing requirements (e.g., in sensor and actuator control applications). A widely accepted body of theory is available for analysis of periodic processing with mandatory deadlines – the *rate monotonic* theory. Systems to which rate monotonic analysis may be applied are sometimes considered to be synonymous with hard real-time.

True hard real-time requirements of the type described in the preceding paragraph are less prevalent in large real-time systems than might be supposed. Many computations have some degree of flexibility in their completion. For example, the Aegis “missile loop” is often used as an example of the need for strict real-time periodic “lock-step” operation. In reality, it need not be adhered strictly to a hard deadline schedule, though it certainly averages out to a periodic schedule. In fact, an average is all that is needed – even considering the physics of the problem – so long as certain longer time-scale guidance timeliness bounds are observed. The loop update rate derives from error containment requirements associated with the missile guidance process. In effect, the missile guidance update periodic contains a designed-in safety margin. This means that it is not strictly a hard real-time process.

Note, however, that, even if absolute “hard real-time” deadlines are not needed, there still exist timeliness bounds beyond which deferral of the computation will result in drastic negative consequences for system operation. In the example of the previous paragraph, the missile error bounds referred to will eventually be exceeded if an update is not available for too long a period. For this reason, both component design and overall system design should carefully implement mechanisms for controlling flow of control during system operation. This is particularly true in

the presence of dynamic RM. Component thread design, buffering mechanisms and priority assignment should be done in a manner consistent with system engineering objectives and in recognition of the interaction of the flow of control characteristics of operating systems and communication mechanisms.

Proper utilization of these control mechanisms during component and system design is the key to simultaneously meeting “hard real-time” requirements while providing balanced service for remaining computational tasks.

**3.2.2.10.1.2 Quality-of-Service.** Closely associated with real-time is the concept of Quality-of-Service (QoS). QoS is often associated with network and information transfer characteristics. A common usage relates to the interactions between network communications paths and applications that are designed to seek a defined level of network throughput and/or latency. Sometimes network QoS is expressed in terms of “jitter,” a variation in time of arrival of network packets at a destination. In either case, applications operating within a QoS range will either request a specified QoS from the underlying network infrastructure or they will adapt their performance expectations to the QoS available from the network. This, in turn, implies that the network should be capable of managing and allocating QoS for a variety of users.

In many instances from the commercial world, QoS is a marketable commodity, and network service providers charge for higher QoS levels. In such situations, cost becomes the metric whereby user decisions are made concerning (1) what QoS to request from the provider, (2) what adaptations to make in QoS utilization to remain within defined cost targets. Likewise, the provider, motivated by profit, may decide how much aggregate resource to make available for use, based upon market projections and other business considerations.

Parenthetically, it should be noted that such systems generally serve a variety of customers for whom inter-application interactions and dependencies are negligible. This is significant for the following discussion concerning real-time mission critical systems such as Navy combat systems – for which strong inter-component dependencies do exist.

While this guidance document incorporates the network QoS concepts discussed above, it generally employs the QoS idea in a somewhat broader manner. Specifically, the idea of QoS is expanded to incorporate processor resource use in addition to network use. This is a natural extension to the QoS concept for large distributed real-time systems. In the latter case, the components of the system should interoperate, frequently in a strongly interdependent manner. Thus, the activities of each component are potentially dependent on the successful operation of many other components.

In this situation, the concept of components negotiating QoS from an independent market-like service provider has little relevance. Instead, the system – and all of its components – should use the available resource base in a harmonious and cooperative manner. Furthermore, both processing and information transfer assets are clearly contributors to the overall successful operation of the system. This difference becomes even more acute when dynamic RM is introduced into the equation (see Section 3.2.3.8). In the dynamic RM case, both processor and network metrics are defined for measurement and management.

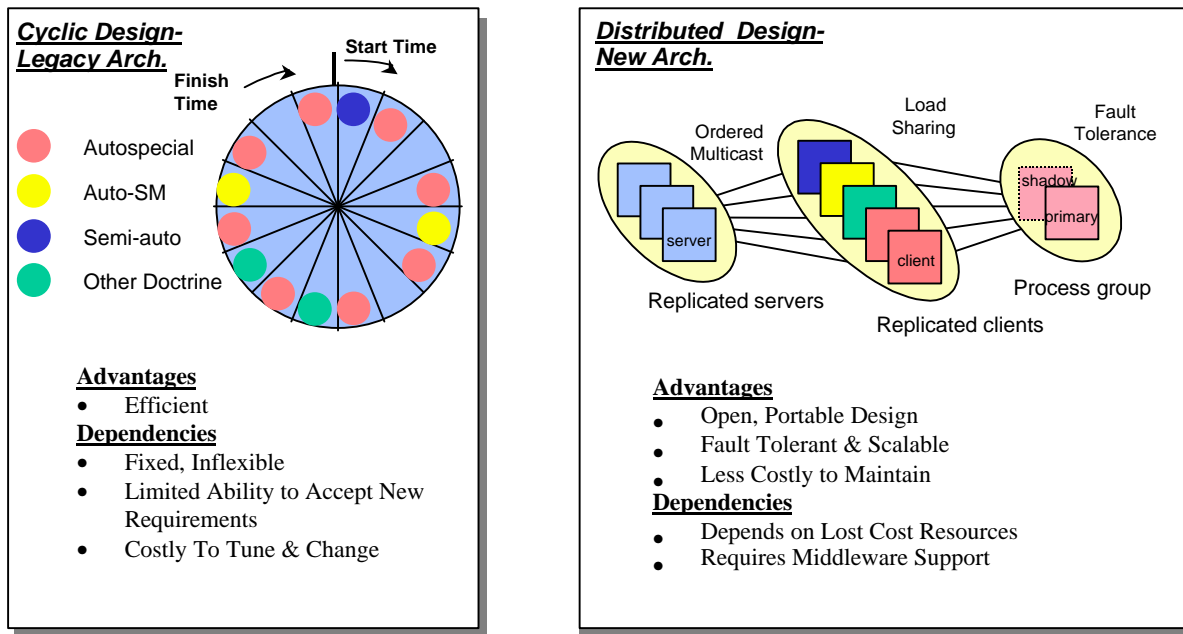
**3.2.2.10.1.3 Distributed Real-time Systems.** Perhaps the most interesting and challenging case of all is the situation faced by designers of large complex mission-critical distributed real-time systems.. For such systems, both real-time scheduling considerations and factors associated with QoS (see discussion below) are present simultaneously and sometimes in the same application. In addition, there are inevitably a substantial number of end-to-end requirements that transit several real-time programs during their accomplishment. These paths often implement low latency responses to critical external events.

Furthermore, while discussions of real-time more often focus on periodic and operating-system scheduling, distributed real-time system designers should inevitably consider as co-equals the timing characteristics of the information transfer mechanisms, as well as those of the operating-system scheduler.

Another consideration is the relationship between periodic processing, event-based processing, and operator control. In large distributed real-time systems; it is difficult to enforce the structuring mechanisms (such as rate monotonic theory) that provide the theoretical foundation for guaranteed hard real-time performance. It is not so much that those structuring mechanisms are not achievable but rather that implementation of them would inevitably engender a greater degree of rigidity to the overall system design than is consistent with modern design practices. Hence, the real challenge for distributed real-time system designers is to ensure that real-time properties are preserved while designing an architecture and a set of components that meet the supportability metrics introduced in Section 1.3.2.

Contemporaneously with the emergence of the need to deal with large complex systems is the development of a number of enabling technologies that were not necessarily invented with real-time requirements in mind. These technologies include networks, distribution middleware and dynamic RM. The composability and supportability benefits of these technologies are compelling; however their suitability for real-time use has not been broadly tested until recently. Fortunately, the technology base introduced in the preceding paragraph has evolved substantially, and in many cases it is now fully real-time capable in all but the most stringent applications.

What emerges from this confluence of design objectives, commercial practice and technology tools is a new system architectural concept that provides a far more flexible basis for meeting supportability metrics. Rather than using shared memory and fine-grained operating-system scheduling mechanisms, this approach connects large numbers of small components via networks and middleware-based inter-component communication mechanisms. The new architecture, contrasted with traditional real-time system design in Figure 9, provides increased flexibility. While it may seem profligate in its use of processing resources, the new architecture trades expanded use of processors and communication, which has become relatively cheap in today's environment, for design flexibility and supportability benefits.



**Figure 9. Real-time System Comparison**

### 3.2.2.10.2 Guidance

Support for meeting this guidance consists of managing a number of computer resources in accord with the nature of the real-time problem. In the guidance stated below, an execution path is a set of processes that should run in succession to perform a complex application function. A path is initiated by an event, and the processes that occur during path execution may result in events that initiate other paths. Initiation of a particular path may be conditioned on completion of one or more other paths. A mission critical path contributes directly to mission success. Failure to meet timeliness requirements for a mission critical path can result in mission failure.

1. Process and task/thread priorities should be used when necessary to ensure that mission critical hard real-time deadlines are met, if necessary, at the sacrifice of soft real-time deadlines.
2. Management of input/output and throughput processing should be accomplished by means of queues that buffer program units from each other and prevent busy states in one part of the system from impacting other parts of the system. Output processing may be handled through a queue mechanism or through the use of separate threads servicing different consumers. Also, most commercial middleware products that support threads (e.g. Network Data Distribution System (NDDS), Object Request Brokers (ORBs)) will handle the output queuing in a satisfactory manner without blocking the application and without application intervention. In the event that a lower level communications mechanism is used (e.g. group communications, TCP, UDP), then queuing of outputs and an associated thread to send the data is highly desirable.)
3. The previous requirement may cause significant performance degradation in synchronous interactions such as a CORBA 2-way invocation, since the thread is blocked until the application gets around to responding. Exceptions to the previous requirement include

three cases: (1) when processing requirements for input is small (e.g., updating a data item in an internal data store); (2) processing of high urgency inputs; and (3) synchronous interactions such as remote procedure call (RPC) style request/reply interactions.

4. Mission critical execution paths in applications hosted by the OACE should be identified and described in terms of process sequence and timeliness requirements (e.g., hard vs. soft real-time, time in which process should complete). Timeliness requirements should be specified for each process in the path and for the path as a whole. (Note: The timeliness requirements will be employed by the RM capability (see Section 3.2.3.8) to predict and manage replication, load leveling or component migration.)

### **3.2.2.11 Process, Thread and Memory Management**

This section provides rationale and guidance for designing efficient real-time multithreaded OA components that ensure correct operation. The section discusses techniques for ensuring that OA components are designed to avoid, as required, the inefficiencies and unpredictability of run-time memory allocation, thread creation and garbage collection. Also, it provides guidelines to ensure that processes and threads are not hard mapped to specific CPUs in processors with more than one CPU. Generally, it is desirable to avoid mapping of threads to specific CPUs; however, this may not be feasible in all cases (e.g., due to poor cache performance or to improve I/O performance). Use of such mapping should be justified by the designer and approved as an exception when necessary. Finally, data protection and access synchronization are part of the overall requirement to preserve data integrity in mission critical and safety critical systems.

#### **3.2.2.11.1 Description**

With modern languages and operating systems, it is possible to defer binding of many system resources until runtime. This has the advantage of flexibility of operation, but it also has the disadvantage of potential inefficiencies and lack of predictability during operations. Since efficiency and predictability are extremely important aspects of most real-time systems, pre-allocation of computing resources such as memory, processes, threads, file descriptors, etc. is strongly recommended. The preferred design approach to control these potential sources of non-determinism is through the use of pre-allocated pools of memory and threads.

Memory allocation is a good example. Programs can allocate memory from a global heap at runtime. Such programs can also release memory during operations. Over time, the global memory heap is almost certain to become fragmented or exhausted unless mechanisms are provided that enable the operating system to reclaim and reorganize unused memory. If not carefully planned as a part of the real-time flow of processing, this activity, called “garbage collection,” can adversely impact the operation of a real-time system.

For a variety of reasons, real-time computer programs are often multithreaded. Threads may be defined to support periodic, event handling, communication with other programs and devices, etc. In multiprocessor systems, it is often possible to constrain individual threads to execute on specific processors. Such constraints usually make programs dependent upon the

implementation details of the specific computer – which in turn makes them difficult to maintain and port to other processors.

In reliable systems, access to data shared among processes and threads should be synchronized and sequenced in order to preserve data integrity, e.g. to prevent corruption of the shared data. To achieve this, certain standard mechanisms should be used.

### **3.2.2.11.2 Guidance**

The scheme most often used for accommodation of the run time range of processing requirements of real-time systems is pre-allocation of pools of resources. In addition, programmer-supplied services are needed to explicitly manage withdrawal from and return to the pool. Process, thread and memory management guidance is as follows:

1. To promote portability and flexibility of design, tactical component threads should not be mapped to a specific processor. There is one exception to this: Real-time programs are usually I/O intensive. Mapping I/O interrupts and processing to a processor dedicated to handling the I/O stream can lead to more efficient use of the remaining processors and is therefore not excluded.
2. Thread safe methods should be used for synchronization of shared data and access to non-reentrant service routines by multiple threads. Such mechanisms include low-level operating system semaphores, adaptive middleware or framework thread-safe libraries, and language-based run-time support structures that provide inherent synchronization semantics, e.g. Ada and Java. Where high-level language constructs support such synchronization with sufficiently low latency, the high-level language constructs should be used.
3. When possible, object-encapsulation mechanisms and thread-safe access mechanisms should be used to encapsulate data that is not needed external to the thread (i.e., such data should be encapsulated within the visibility of the thread).
4. Thread priorities should be established in order to ensure that urgent processing, for which defined latency requirements exist, are able to preempt processing of lower importance, except in cases where state data consistency would be affected negatively by the use of thread priorities.
5. For mission critical components, process priority distinctions should not be required to ensure correct system operation and satisfactory performance. Operating system scheduling management of execution flow should be implemented by means of thread priorities. Thread scheduling considerations are the primary determiner of real-time performance, even for situations where multiple components are allocated to the same processor. Process priorities, if used at all, should be confined to setting non-mission critical process priorities to a low level so that such processes do not interfere with mission critical components.
6. Mission critical applications should not allocate memory or threads at run time, other than as an initialization activity.
7. Whenever possible, applications should actively managed system resources such as memory allocation and reallocation, even if this management should be performed in application code. Management of dynamically allocated assets such as garbage collection should be entrusted to system software only where system mechanisms exist

by which applications can control the flow of processing associated with the system's management of those run time allocated resource. Such mechanisms should be used by the application so as to ensure that system dynamic asset management operations do not interfere with the operation of mission critical components.

8. Components should maintain precise control of the accretion, use and disposition of system resources (memory, threads, devices, system services, etc.), so that at shutdown time they can either return such resources to the operating system or a system-wide resource control capability can ensure that the resources are released cleanly.
9. All memory and thread management should be performed in a thread-safe manner (i.e., access to data structures and non-reentrant service routines that may be shared among multiple threads should be protected by thread-safe mechanisms as discussed in Item 3 above.)

### **3.2.2.12 Data Brokers**

This section provides rationale and requirements for providing data broker and/or layered adaptation mechanisms to accommodate interfacing to legacy capabilities.

#### **3.2.2.12.1 Description**

Data brokers are translator components that are positioned between two disparate systems to facilitate communications. They adapt the electrical interface, protocols, formats, and syntax of one system to that of another. Since the data broker itself is software that should be developed and maintained, this solution leads to increased development and maintenance cost, albeit likely a small increase. Furthermore, data brokers are only useful in situations when performance objectives are not adversely impacted by the overhead that the data broker adds. Thus, a data broker is a compromise solution and it should only be used when large-scale modifications to either of the interfaced systems are prohibitive due to cost or schedule.

Data brokers provide a standardized, consistent API for legacy applications that, when employed, "translate" the data representation of messages into one that is compliant with the OACE architecture. Conversely, information that is sent to the legacy system from the OACE architecture is translated into a data representation that is compatible with the legacy system. Data brokers permit the phased implementation of the OACE architecture by insulating legacy applications. Furthermore, when properly designed, data brokers ensure that changes will not be required in the OACE components if and when the legacy components are modified to become OACE compliant.

Adaptive layers, generally provided as linkable libraries, confer similar isolating services. The adaptive layer exports an interface that hides the details of the legacy system to which communications should be maintained.

The use of data brokers and adaptive layers will contribute to reduce efforts in maintenance and upgrades by isolating the legacy applications from OACE.



#### **3.2.2.12.2 Guidance**

The guidance for legacy capture within OA is that a data broker and/or adaptive layer approach may be used to interface with legacy components..

### **3.2.3 Open Architecture Computing Environment**

The technology base, or infrastructure, for OA is referred to in this document as OACE. Since tactical applications will be implemented as a set of loosely coupled software components, the fundamental requirement for OACE is to provide a distributed real-time computing environment for NWS elements – and possibly ACS elements in the future. For these elements, issues of fault tolerance, timeliness, scalability, expandability, and maintainability are vital.

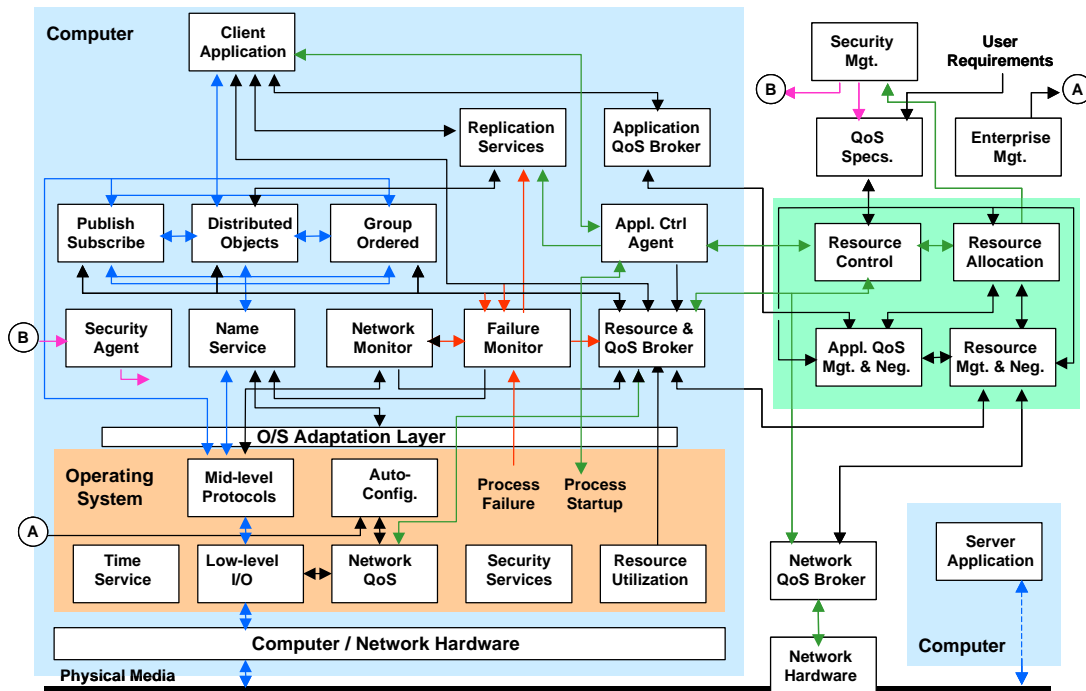
OACE is a high-performance computing environment for distributed processing, and it is capable of controlling and executing distributed tactical applications. OACE provides a common computing environment with services for application loading, resource management, inter-application communication, fault tolerance, and adherence to real-time application requirements. The resource base controlled by OACE consists of a set of computers, internal and external network interconnection equipment, network media, operating and control software, communication software, and interface software. Distributed application software will load and execute in the OACE environment.

Listed below are a number of key technical attributes of the OACE:

- CORBA inter-component distributed object mechanisms when performance permits
- Distribution of continuously refreshed data via publish/subscribe protocols
- Maintenance of state data coherency via group/ordering protocols
- Dynamic resource management, with allocation of portable processing components (i.e., “run-alike” across brands, operating systems, etc.) onto a pool of virtually homogeneous computers.
- Instrumentation data from operating systems, network components and applications
- Integrated failure management services across infrastructure components
- Information security integration with RM
- Processors with suitable responsiveness and queuing of I/O interrupts
- Effectively uniform switch interconnect with QoS features, e.g. the network Resource Reservation Protocol (RSVP)
- Time synchronization via system-wide service (e.g. Network Time Protocol (NTP))
- Open, standards-based COTS products (as near mainstream as possible)

The architecture of the OACE supports the principles, structure, and components that form the basis upon which the application components of OA are built. Thus, the attributes of the OACE support development of application components in accord with the guidance contained in Section 3.2.2. It includes technologies needed to provide guaranteed QoS to real-time applications. Figure 10 illustrates the functions and interrelationships of OACE.

**Figure 10. Open Architecture Computing Environment**



Note that this representation is given in functional form and does not necessarily represent a physical design. For example, three classes of communication protocols are shown: distributed objects, publish-subscribe and group-ordered communications. These classes exist as separate tools at present, forcing designers to select one and only one of the three for a particular interface; however this may not be true in future products. Current middleware marketplace evolution appears to be in the direction of combined functionality in a single product (e.g., a fault-tolerant CORBA based upon group communications with publish-subscribe functionality).

This architecture provides the capabilities and infrastructure needed to construct multi-component, replicated, distributed object real-time systems. The structure is repeated throughout the computers of OA. Thus, although the diagram elaborates the technology base for a computer containing a client component, the computer hosting the server application contains a comparable structure. Also executing within the distributed system's collection of computers is a set of RM components that interact with the computers, network components and applications of the distributed system to provide QoS management.

The OACE consists of the layers listed below:

- Cabling and cabinets
- Information Transfer
- Computing Resources
- Operating Systems
- Adaptive Middleware
- Distribution Middleware

- Middleware Frameworks
- Resource Management
- Tactical Support Services

Each layer constitutes a part of the underlying technology base and is discussed in the sections that follow.

The OACE provides a critical set of functions required for the proper operation of applications components within OA. Among these characteristics is a capability for the OACE to monitor itself. Another is portability of the OACE technology base across a variety of COTS products and standards from which future technology refreshes are likely to be chosen. COTS hardware and support software is expected to evolve rapidly. Therefore, the OACE should be designed to support cost effective modification, addition, and replacement of hardware and software components.

### **3.2.3.1 Cabling and Cabinets**

This section provides rationale and guidance for the cabling plant and cabinet architecture of OA. The computing resources should be accommodated within the footprint and connectivity constraints of existing cabling and cabinetry.

#### **3.2.3.1.1 Description**

One of the key boundary conditions for development of OA is that it should be deployable in pre-existing ships without impacting the cable plant and cabinetry already installed on those ships. This condition should be observed to the maximum extent possible. However, future combat system evolutions may necessitate revisiting this requirement, independent of OA constraints.

A OA re-implementation of the tactical programs will almost certainly fit inside the current cabinet footprint, with minimal effect on legacy I/O nodes. They will all become I/O processor gateways,. However, the network switching design is likely to be quite different, due to the presence of more computing nodes. Additional study will be required to determine if the installed base of fiber can handle it, despite its relatively flexible patch-panel-based design. For upgrades starting with NTW Block I and progressing to a potential future SPY-1E, the footprint issue would have to be re-examined.

#### **3.2.3.1.2 Guidance**

Guidance for use of existing cabinet footprint and cabling is as follows:

1. Unless compelling reasons exist for major change, OA should maximize compatibility with the cable plant design on existing ships.
2. OA should be designed to facilitate installation of its components in equipment enclosures that meet space and weight budgets for legacy ships. Physical design constraints should apply to cabinets, computer subassemblies, junction boxes, network

media, and all other items comprising the physical manifestation of OA.

### **3.2.3.2 Information Transfer**

This section provides rationale and requirements for designing the OA information transfer capability. The goal is to ensure that information transfer bandwidth and latency requirements are met and that no single location within the network represents a constraint on information flow or process to processor allocation. This latter feature is intended to support the allocation and reallocation of processing tasks at run-time.

#### **3.2.3.2.1 Description**

The network will provide the primary means of information transfer within OA. While a limited number of point-to-point interfaces will continue to exist for some time, NWS and ACS elements will continue to migrate toward a full-network implementation.

**3.2.3.2.1.1 Basic Network Functionality.** Three basic functions provided in the information transfer infrastructure are connecting a processor to the network; interconnecting all the components on a single network; and connecting networks together. In an OA, these functions are to be provided by (1) network interface cards (NICs), (2) switches, and (3) routers respectively.

A NIC is the means by which a processor accesses the network to send or receive data. A NIC is required for each type of networking technology used by a processor, e.g., Fast Ethernet, Gigabit Ethernet, or Asynchronous Transfer Mode ATM.

A switch is a networking device that receives and transmits data on a number of ports. Unlike shared network architecture, a switched architecture dedicates each port on a switch to the device connected to it. In other words, a processor connected to a Fast Ethernet switch has a dedicated 100 Mbps path to transmit and receive data. The internal components of the switch temporarily connect switch ports in order to enable transfer of data to other computers or to additional switch components. The number of active ports a switch is able to realistically support depends upon the throughput of the backplane of the switch.

Routers provide for the forwarding of data across different networks (a network is defined here as a single broadcast domain). Routing functionality may physically be provided in switch hardware, but the functionality is distinct from an architectural perspective.

**3.2.3.2.1.2 Communication Mechanisms.** Three basic communication mechanisms exist in the information transfer infrastructure: Uni-Cast, multicast and broadcast. Uni-Cast messaging involves a sender transmitting a message to a single receiver.

Both multicast and broadcast involve a sender transmitting to multiple receivers. The functionality of these two mechanisms is distinguishable at two layers. At the higher Internet Protocol layer, broadcast messages are transmitted to all receivers within the defined domain. For multicast, a sender transmits to a selected group of receivers defined by membership in a multicast group.

The way in which the higher layer message is instantiated on the lower layer depends upon the type of network media being utilized (e.g. ATM, Ethernet). Broadcast and multicast messages from the higher layer are transmitted as broadcast messages at the lower layer in a shared media. Techniques such as Internet Group Management Protocol (IGMP) snooping have evolved in shared media to filter multicast messages to just those members of the multicast group. When one or more senders should transmit the same message to multiple receivers, multicasting is more efficient than point-to-point uni-casting.

**3.2.3.2.1.3 Additional Network Features.** The network should also support one of Open Architecture's most important design characteristics: deferred binding of program components to physical resources. Middleware supports deferred bindings by providing location transparent communication services. RM supports deferred bindings through a capability to assign computational components to physical computers just prior to operations, or even dynamically, as the system runs.

Another key requirement is that the network should provide uniform service performance to real-time applications. Dynamic RM (DRM) enables much more flexible binding of computing tasks to processors. Potentially, DRM can allocate any program capable of running on the real-time computing unit (see Section 3.2.3.3) to any computer of this class. The corollary to this allocation capability is the fact that any computer within this class may eventually originate a message destined for any other computer of like class in the system.

For this reason, all paths through the network connecting real-time units of computing resources should be effectively uniform in performance, or at least they should provide a specified minimum performance. In other words, each network path should be equally capable of supporting the overall required latency and throughput requirements for inter-computer message traffic between mission critical components.

Switched networks are most likely to possess this virtually uniform performance characteristic. For such networks, the uniformity of performance requirement is generally only a concern on the individual switch-to-switch links. By definition, traffic on edge-device-to-switch links is limited to the available link bandwidth, which is, in turn, entirely devoted to the single device to which it is attached. Some configuration control over internal network characteristics is, therefore, likely to be needed. This control, both in terms of instrumentation and reconfiguration, should be entrusted to the RM capability.

#### **3.2.3.2.2 Guidance**

The following constitutes guidance required of the information transfer capability by OA.

1. The information transfer infrastructure should provide an effectively uniform information transfer capability from any source in the network to any destination in the network with respect to connectivity, latency and throughput. Effectively uniform means that all network paths will provide at least a defined minimum level of service.
2. The information transfer infrastructure should provide fault tolerance of information transfer paths.
3. The information transfer infrastructure should provide standards-based QoS control

- mechanisms for bandwidth reservation (e.g., RSVP).
4. The network infrastructure should provide standards-based mechanisms to dynamically modify routing paths for the purpose of ensuring the uniformity requirement following damage or disruption to the overall network. Conversely, such standards-based mechanisms may be used to insure maximum QoS to components requiring it.
  5. A common set of information transfer services should be provided to coordinate and facilitate internal communications. OACE should employ industry- standard protocols, such as the Internet protocol suite (see Section 3.2.7 for guidance on standards selection).

### **3.2.3.3 Computing Resources**

This section provides rationale and guidance for selection of computing resources for OA from the market of high-performance, low-cost commercial processors.

#### **3.2.3.3.1 Description**

From a maintainability standpoint, one of the benefits of distributed computing is the capability to modify system capacity without requiring changes to applications. However, due to the rapid turnover of processor technology, this necessitates portability of applications across various commercial processors. As a means for reducing costs, the Navy intends for OA to take advantage of, as much as possible, the high-performance low-cost mass-produced commercial computing market, while maintaining this portability condition. In order to do this most effectively, processor selections for OA should be made from mainstream single- and dual-processor boards. Niche market products such as many-processor symmetric multiprocessors should be avoided.

Real-time systems are almost universally I/O intensive systems. Handling many I/O interrupts per second can impose a large overhead penalty on a single processor computer. For this reason, some implementations utilize a second processor to field interrupts and perform initial I/O processing. This is a critical factor in equipment selection.

There are multiple ways to handle this; one method is to use symmetric multiprocessors and the operating system's inherent support for multiprocessing. In this scheme, I/O interrupts may be bound to one processor exclusively. In any case, the critical design goal is to shield at least one processor from the burden of I/O handling, so it can make progress on application computations.

The Naval Warfare Systems included in OA contain considerable breadth of processing requirements. While it is certainly possible to meet all OA processing requirements with one class of processor, it may prove to be advantageous to provide multiple classes of processing capability within the scope of OA. The guidance here is to move away from symmetric multiprocessors; however, a multi-tier computing resource approach need not exclude computer/operating system combinations that provide less than the most stringent level of real-time capability and performance. For example, the processing resources devoted to display and other soft real-time activities need not meet the challenging requirements imposed by a radar control program or signal processor. In addition, it should be noted that where processing requirements dictate a specialized solution, such solutions might be adopted as exceptions to the

guidance below. This partition of resources will reduce but not eliminate the flexibility inherent in the dynamic RM capability.

### 3.2.3.3.2 Guidance

OACE computing resources should consist of hardware and system software with sufficient capacity and performance to host the distributed OA applications.

1. A *virtually homogeneous* environment should be provided for OACE compliant components within the OA. The term, virtual homogeneity refers to the characteristic of a system whereby each computing element performs similarly enough so that it meets the minimum performance, latency, and determinism requirements of the components that make up the system. Virtual homogeneity allows realization of the benefits of commonality without requiring absolute equipment commonality.
2. Processors with direct IO connections, and therefore fixed computational task-to-processor mapping, are exempt from the virtual homogeneity guidance, item 1 above.
3. The OACE computing architecture should be based upon fundamental uniform building blocks called *units of computing*. The unit of computing is defined as a combination of processing speed, memory, network performance, latency, and a selection of industry standard interfaces. The unit of computing concept provides a hardware modularity that permits the development of a system that is scalable, fault tolerant, allows for resource sharing, and enhances the portability of the functional design.
4. Units of computing should be combined to accommodate, via simple scaling, a variety of applications that would otherwise require unique processing solutions.
5. Each unit of computing should consist of a single or dual CPU under the control of a single instantiation of the operating system. The unit of computing should be sized to meet the requirements of the distributed tactical applications and should be characterized by the following performance parameters. Details follow below in Table 4. Units of Computing Performance Parameters.
6. The values selected for parameters in Table 4. Units of Computing Performance Parameters should be sufficient to accommodate the needs of the tactical applications and should be representative of low-cost high-performance mainstream commercial computers. The values for each of these parameters should be selected, such that the units of computing are not excessively oversized for the majority of requirements, yet can be scaled to support the needs of computing intensive combat system elements.
7. Scalability of the OACE may be achieved by use of the following mechanisms:
  - 1) Additional units of computing
  - 2) Incremental upgrades of components in the unit of computing (e.g., memory)
  - 3) Technology refresh

**Table 4. Units of Computing Performance Parameters**

Parameter	Description
SPECint95	SPECint95 and SPECfp95 are industry standard measures of processing power.
SPECfp95	

Parameter	Description
Memory	The quantity and type of memory for a unit of computing should be specified.
High Speed Network Interface	The type of high-speed network interface should be specified to facilitate combat system wide integration.
Transmit I/O Capacity	These parameters ensure sufficient bandwidth and determinism between distributed components.
Transmit I/O Latency	
Interrupt Response Latency	Interrupt response latency is the maximum time (in $\mu\text{sec}$ ) allowed from the generation of a hardware interrupt by a device to the execution of device driver or kernel code servicing the interrupt. This worst-case time includes all possible periods when interrupts are disabled in order to guard some operating system critical section.
Process Dispatch Latency	Process dispatch latency is the maximum time (in $\mu\text{sec}$ ) allowed from the generation of a hardware interrupt by a device to the execution of user code requesting service from that device. This worst case time includes the interrupt response latency, the time required to execute device driver or kernel code servicing the interrupt, overhead required in context switching to the user process, and any critical sections when context switching is disabled by the operating system.
Fast Clock Resolution	This is the period (in $\mu\text{sec}$ ) of the fast clock
Fast Clock Access Time	This is the maximum time (in $\mu\text{sec}$ ) required for a user process to read the fast clock. A local clock should maintain time values that are always monotonically increasing such that time does not run backward. In addition, each access to the clock should produce a unique value such that no two consecutive clock accesses will result in the same time value being returned.

### 3.2.3.4 Peripherals

This section provides rationale and requirements for selection of computer peripherals, special purpose devices and other add-on computing capabilities. The guidance is as follows:

1. Interfaces to commonly available commercial information technology peripherals should utilize commercial mainstream industry standards.
2. Special purpose peripherals, such as NTDS interfaces, IRIG time boards, etc. should use interface protocols appropriate to the device type. Every effort should be made to provide a portable and maintainable interface mechanism.



### **3.2.3.5 Operating Systems**

This section provides rationale and guidance for selection of operating systems in general and POSIX compliant operating systems with suitable real-time features for mission-critical applications.

#### **3.2.3.5.1 Description**

In the operating system domain, a number of features are required. The operating system should support a system of priorities to ensure those real-time processes and threads are dispatched when ready to run. Support should be provided for kernel pre-emption in non-critical region processing. The operating system should also provide mechanisms to control priority inversion (e.g., priority inheritance protocols). Network drivers and other I/O operations should fall within the scope of the operating system real-time priority structure, and they should provide bounded delivery latencies under defined circumstances.

#### **3.2.3.5.2 Guidance**

Guidance for operating systems is as follows:

1. The operating system should provide standardized APIs. (The POSIX standards provide one mechanism for attaining this.)
2. In cases when use of POSIX-based operating system APIs is not possible, the operating system services should be provided through a portable standard operating systems adaptation layer.
3. All operating systems intended for use in mission critical applications and support software should support the completion of hard and soft real-time deadlines required by the distributed tactical components
4. For distributed components with real-time processing requirements, the operating system should provide standardized real-time scheduling facilities.
5. All operating systems should provide mechanisms for assigning application level priorities to processes and threads, for kernel pre-emption and re-entrance and for protection against priority inversions. Kernel critical regions should be bounded with values not to exceed those established by application requirement analysis.
6. Support services include operating system services, file management services, standard time services, and other utilities required for proper system operation, including programming language facilities.
7. Software for support services should make use of industry standards and products to the maximum extent possible. Standardized file access mechanisms, particularly for shared network access such as the Network File System (NFS) should be provided as part of the support services.
8. At a minimum, operating systems should support system scope threads.

### **3.2.3.6 Adaptation Middleware**

This section provides rationale and guidance for selection of adaptation middleware and products that export a POSIX-like operating system interface to application components.

Middleware may also be used to export a common distribution middleware interface (see Section 3.2.3.7).

#### **3.2.3.6.1 Description**

Adaptation middleware products export a standardized interface upward (i.e., from multiple-operating systems to a single API). The Adaptive Computing Environment (ACE) from Washington University of St. Louis, MO. provides adaptation layers for most commercially available operating systems. ACE exports a standard set of APIs and services to applications. As such, it strongly promotes application portability. Other adaptation layers may be developed for middleware functions. These layers often provide a way of making non-standard products appear to be standards compliant. In this way, application portability is enhanced.

#### **3.2.3.6.2 Guidance**

The guidance for adaptation middleware products is that a widely available and commercially supported adaptation layer such as ACE should be considered for utilization within OA to maintain application portability across operating systems.

### **3.2.3.7 Distribution Middleware**

This section provides rationale and guidance for use of distribution middleware in OA implementation.

#### **3.2.3.7.1 Description**

The term “middleware” is most often used to describe support software that facilitates interactions between major software components and masks differences in language, platform characteristics, message formats, communication protocols, data structures, and other factors. A number of communication middleware protocols have value in designing large, complex real-time systems. Furthermore, an evolution in recent years has occurred with respect to middleware semantics, driven in part by the widespread acceptance of the object-oriented paradigm among developers. Three classes of distribution middleware that provide capabilities relevant to OA interprocess communication functionality are as follows:

- Distributed object computing
- Publish/subscribe data distribution
- Group-ordered communication

**3.2.3.7.1.1 Distributed Object Computing.** Within the domain of distributed processing, the distributed object computing model has emerged as a robust, standards-based technology area that is ready for use in real-time systems.

At present, three major distributed object models exist: Microsoft’s Distributed Component Object Model (DCOM), Java’s Remote Method Invocation (RMI) and Enterprise Java Beans, and the Object Management Group’s (OMG) CORBA standard. Of these three,

CORBA offers the widest range of competitive vendor alternatives for a multi-language, multi-platform environment. In addition, OMG has recently released a real-time CORBA specification and several conforming implementations are available.

CORBA provides location transparent mechanisms to create object-oriented interfaces between distributed components. When designing a distributed-object model, one should consider the impact of the design on overall system performance. While creating network objects is rendered straightforward by CORBA, performance still depends on the network communication infrastructure.

System partitioning into numerous small objects provides a means of more precisely modeling the physical world; this point is reflected in the component partitioning section above. However, if carried to an extreme, a design predicated on this model may overload the network infrastructure. Consequently, achieving a balance is crucial between defining modestly sized functionally homogeneous components and defining components that are too small to operate efficiently.

**3.2.3.7.1.2 Publish/subscribe data distribution.** Within the domain of the publish/subscribe paradigm, messages are published to a multicast group for which membership is encapsulated by middleware. To become consumers of messages, clients register with the middleware using group names indicative of the type of data required. Servers publish data using the appropriate group names. The middleware uses the group names to ensure that data is routed to the appropriate clients, thus divorcing the server of knowledge regarding which clients have subscribed for data updates.

The publish/subscribe paradigm is well suited for location transparent distribution of continuously refreshed data such as distribution of track data to consumers as described above. In general, publish/subscribe performance tends to be better than other middleware approaches and, as such is often the first choice for high data rate applications.

The publish/subscribe paradigm does have limitations that need to be accounted for when designing a distributed system. One limitation is the lack of sufficient functionality to fully support the replication mechanisms required to implement fault tolerance design patterns (see Section 3.2.2.8). This limitation can be overcome through the use of group-ordered communication protocols (see Section 3.2.3.7.1.3). Another limitation is the lack of a distributed object interface, such as that provided by CORBA.

At present, there is no standard for publish-subscribe functionality. However, several readily available commercial publish-subscribe products provide similar services. Furthermore, OMG has instituted a publish-subscribe-working group. That group recently issued a Data Distribution Request for Proposals (DDRFP). It is believed that the output of this effort will eventually be a standard for distributed object publish-subscribe functionality. In the interim, the best approach is probably development of an adaptive layer for publish-subscribe protocols. Such a layer would enable swapping out publish-subscribe products over time with a minimum of disruption to application code.

**3.2.3.7.1.3 Group-ordered communications.** Group-ordered communications

middleware provides support for sophisticated reliable inter-process communications in distributed systems. Central to group communications are the process group semantics. In many respects, a process group is similar to a multicast group:

- A process group is a set of communication endpoints.
- Messages addressed to a process group are distributed to all members.
- Distributed applications join a process group to exchange messages.
- Multiple groups can be created within a distributed system.
- Any application may simultaneously be a member of multiple process groups.

Essential group membership semantics differentiate a process group from a simple multicast group:

- Each member of the process group has access to specific information about the other group members.
- A group member can determine how many members are in the group, and perhaps even which specific processes executing on other nodes in the distributed system are members of the group.
- Process group members can receive events indicating that a new member has joined the group or a current member has left the group.

These additional semantics support the development of application redundancy for dependability, fault tolerance, and load sharing. Additional features such as atomic multicast (also known as safe multicast) ordered messages, group-view changes that are ordered with respect to the message stream, and ordered state transfer provide significant benefit to the developers of reliable, distributed systems.

Approved standards for group communications tools do not yet exist. The Institute for Electrical and Electronics Engineers (IEEE) POSIX 1003.21 Standards Committee is addressing membership-aware multicast group communications APIs, but it is not addressing support for higher-level group communications features, such as message atomicity and ordering.

The work in fault-tolerant CORBA benefits greatly from a standardized group communication infrastructure. The demand for fault tolerant CORBA is growing, and the OMG Real-time Special Interest Group (SIG) has published a suitable specification. Development of fault-tolerant CORBA implementations, already available in limited numbers, will have a positive influence on product availability in this critical area.

Another significant middleware consideration that contributes to successful construction of open systems is data marshalling and endian conversions. CORBA provides a common data representation and data marshalling. Some publish/subscribe tools do, as well, although seldom do they provide efficient mechanisms. Many do not provide such representations and endian conversions. A Receiver-Converts strategy, therefore should be adopted for all messages, as well as common data representations, including 8, 16, 32, and 64 bit signed and unsigned integers, floating point, and Booleans. Message data should be aligned on word boundaries.

### 3.2.3.7.2 Guidance

Guidance for distribution middleware and its use within the OA are as follows:

1. Distribution middleware should be provided to support the interactions among the distributed components, CI/NDI support services, standard services, RM services, and legacy systems (via data brokers) that make up OA (see Figure 4).
2. Distribution middleware software should be composed of COTS/NDI components when performance requirements permit. (Note: domain specific components may be required to perform functions not included in distribution and adaptation middleware, [e.g., NTDS interfaces]. These functions may be generically referred to as middleware, but since their counterparts do not exist in the commercial world, they are exempt from the COTS/NDI constraint.)
3. When commercial distribution middleware standards do not exist, appropriate adaptation layers should be developed that implement an interface representative of extant commercial products.
4. Middleware technology for inter-component communication interfaces should be chosen in accord with the requirements of the interface. Whenever performance requirements and component design pattern permits, the use of a distributed object interface should be used. The preferred distributed object standard is CORBA. High-volume distribution of continuously refreshed data, such as tracks or navigation information, should be accomplished via publish/subscribe methods. Middleware services for group ordered communication should be used where state data coherency and ordering properties are required.
5. Middleware should be used that provides enabling services for communications among heterogeneous platforms. All middleware tools used in the OA baseline should provide data marshalling and endian conversion services. These services should be used where performance requirements permit.
6. Where available and otherwise suitable, middleware should be selected that provides explicit control over classes of service (i.e., QoS characteristics such as timeouts, latency, throughput, message delivery guarantees, etc.). An example of explicit QoS control would be a middleware product with an API allowing applications to specify a bounded latency associated with information transfer or distributed object invocation.
7. Middleware-based QoS features should not be used in ways that create system level resource use conflicts among components. Such conflicts should be ameliorated through use of dynamic RM (see Section 3.2.3.8). Use of RM's QoS control permits resource tradeoffs to be managed explicitly at the system level rather than contained implicitly in the aggregated impact of component-level patterns of QoS invocations.
8. All middleware tools used in OA should provide an API that limits application exposure to operating system and protocol specific functionality, product specific message formats and data structures. Except when required to meet specific performance or functional requirements, middleware tools should be chosen that allow the encapsulation of such functionality, for example, through initialization calls or configuration files.
9. For middleware protocol types where standards are not yet available (e.g., publish/subscribe) a standard middleware adaptation layer should be developed that encompasses functionality typical of each relevant available class of products. This adaptation layer should export a uniform API to applications using such services. Such

layers should be removed when suitable standards and standards-based products become available.

### **3.2.3.8 Design Patterns and Frameworks**

This section provides rationale and guidance for the use of frameworks in the OA product set.

#### **3.2.3.8.1 Description**

According to Gamma, Helm, Johnson and Vlissides in their book *Design Patterns: Elements of Reusable Object-Oriented Software*, a framework is “a set of cooperating classes that make up a reusable design for a specific class of software.” The three key differences between design patterns and frameworks are:

1. Design patterns are more abstract than frameworks. Frameworks can be embodied in code, but only examples of patterns can be embodied in code.
2. Design patterns are smaller architectural elements than frameworks. A typical framework contains several design patterns, but the reverse is never true.
3. Design patterns are less specialized than frameworks. Frameworks always have a particular application domain.

A framework may provide a reusable design for all or merely a part of an application. For example, a fault tolerance framework may provide the structure and interaction of classes that support building a fault tolerant application. Fault tolerance is an issue that is pervasive and a fault tolerant framework would affect the entire application. A tasking framework may provide the set of classes and class interactions that support a particular concurrency model. Although the use of concurrency is extremely important in the component architecture, the model of concurrency can vary across different parts of the component. Care should be exercised during component design to insure that instances of framework use do not conflict with instances where underlying services are invoked directly. The use of frameworks provides structuring mechanism, design patterns and code reuse opportunities that can be leveraged by application developers to reduce development time and life-cycle costs.

#### **3.2.3.8.2 Guidance**

Guidance for the use of frameworks within the OA is that frameworks should be leveraged to facilitate application structure commonality and code reuse where appropriate frameworks exist and have been identified as supplying needed capabilities (e.g. fault tolerance, concurrency, database access, instrumentation).

### **3.2.3.9 Dynamic Resource Management (DRM)**

This section provides rationale and guidance for management of computing resources and dynamic allocation of components to the pool of virtually homogeneous computers that constitute the OA equipment suite.

#### **3.2.3.9.1 Description**

The purpose of RM is to provide run-time control of the computing resources in the system. In providing this capability, two levels of control are possible: static and dynamic. With static control, configurations are fixed at design time and the run-time resource manager uses these configurations to perform startup, health monitoring, reconfiguration, and shutdown.

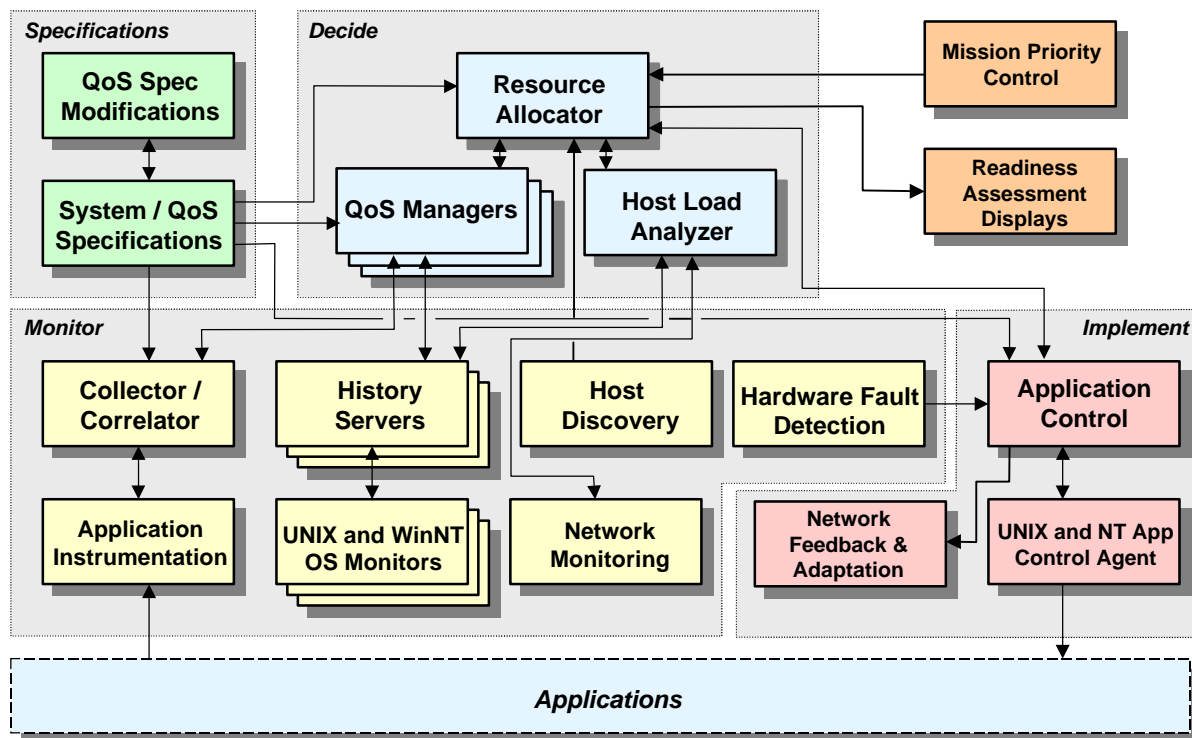
Similar to static control, a dynamic control performs startup, health monitoring, reconfiguration, and shutdown. However, with dynamic control, the resource manager uses application and system performance requirements to derive and execute appropriate system reconfigurations at run-time, in addition to supporting pre-defined configurations. Dynamic reconfiguration of this nature enables the resource manager to ensure that the system continuously operates to provide a defined level of QoS.

Figure 11 provides a RM architecture that features both static and dynamic capabilities.

RM components provide monitoring and control of network, operating system, middleware, and application resources within the distributed computing environment. Resource and QoS management services are intended to provide the ability to dynamically configure and reconfigure hardware, operating system, and application components to ensure that desired or negotiated application, path, and mission-oriented QoS requirements are met at run-time.

To be most useful and least intrusive, the RM capability should introduce little or no impact on basic application design choices. It should not, for instance, arbitrarily force selection of single-threaded vs. multi-threaded applications, periodic processing vs. event driven, blocking vs. non-blocking I/O, or the choice of middleware approach.

The high-level architecture of the RM capability is a feedback control loop consisting of monitoring, decision-making, and control (configuration adjustment) components. System designers specify the required performance of each component. Application monitoring instrumentation measures component performance with respect to specified requirements. When a violation of required performance takes place (or is projected to take place) the RM decision-making capability utilizes status, load, performance, and fault metrics gathered by components that monitor processor and network load. The decision-making capability then determines if, when, and how reconfiguration actions should be carried out. The resource control components then map RM control orders to the low-level control actions that should be carried out to implement the desired reconfiguration actions.



**Figure 11. Resource Management Architecture**

### 3.2.3.9.2 Guidance

RM monitors system performance to assure that the system meets the mission priorities established by ship's command. RM supports fault detection and recovery in cooperation with application components. Guidance for RM, both static and dynamic, are given below:

1. RM should have the capability to determine and control system configuration, load and initialize application components, allocate computing resources, and monitor and control the use of resources.
2. RM should have both an automatic mode of operation and an operator-controlled mode of operation.
3. RM should provide a capability for engineers and operators to statically define and activate system configurations, including processor assignment, communication routing, and other factors. A sufficient number of alternate static configurations should be established to meet availability requirements via automatic- or operator-initiated reconfiguration in response to system faults and loss of resources.
4. RM should have a capability for engineers to define QoS and performance requirements for application components. The QoS specification language should be independent of programming language and application domain. QoS specification capabilities should include the ability to define a number of processing metrics, including response times, periodic processing intervals and duration's, multi-component end-to-end latencies, load-balancing metrics and failure conditions and recovery actions.
5. RM should provide a capability to assign QoS requirements dependent on defined system modes of operation (e.g., non-combat operations vs. combat operations). A



capability should be provided whereby system engineers can define required system configuration reallocations in accord with the new mode of operation.

6. RM should coordinate the recovery actions of all system components (see also Section 3.2.3.11). RM should provide services in support of fault detection and recovery. Faults should be detected in processors, peripherals, network components, middleware and application computer programs.
7. RM should interoperate with component recovery mechanisms; i.e., time critical applications should be designed to perform quick response failure recovery without intervention by RM. In performing the quick response recovery, if a time critical component falls below an acceptable level of redundancy, RM should operate to restore the lost redundancy after the fact, subject to resource availability. Non-time critical applications should be designed so that they may be restarted under RM control.
8. RM should have the capability to continuously monitor mission and system priorities, application performance and the status and utilization of system resources (see application and system resource instrumentation guidance in Section 3.2.3.10). When placed in the automatic mode of operation, RM should allocate resources dynamically at runtime to tasks without operator intervention as required for system optimization, tactical load balancing and fault recovery, or when faults are anticipated.
9. RM should provide a capability to define end-to-end QoS performance requirements across multiple components and to monitor performance, latency and load along these multi-component paths. Upon detection of a violation of the specified QoS parameter, RM should provide a capability to reallocate components for improved performance and to initiate load management via activation of multiple-replicated copies of a component.
10. The dynamic RM capability should have the capability to adjust internal network characteristics (e.g., reserved bandwidth, routing, etc.) in order to ensure tactical performance.
11. RM should provide repeatable and testable mechanisms to ensure closed-loop stability with respect to resource allocation operations.
12. To ensure component compatibility with RM, component developers should design certain features into system components. These features include the following:
  - 1) Specification of application component configuration and performance requirements via a QoS specification language
  - 2) Initial and statically defined system configuration
  - 3) Application and support component startup and shutdown services
  - 4) Operating system level applications monitoring
  - 5) Middleware-level application monitoring
  - 6) Application instrumentation and performance monitoring (see Section 3.2.3.10)
  - 7) Direct RM-to-application control capabilities
13. RM should provide operator display screens for both engineering activities and for deployment operations. Deployment displays should be limited in complexity and supported with automation to reduce required crew workload.
14. RM performance should be sufficient to support all specified OA tactical performance requirements. Detailed RM performance requirements should be defined during OA development.
15. RM should include functionality to monitor and display data internal to OACE components, including RM itself. Human interaction with OACE via RM should

include, but not be limited to, mode selection, operability tests, trouble-shooting, and maintenance. Maintenance actions should include modification of the system configuration following a change to a processor or network component and input of static component allocation data during an engineering maintenance activity.

16. RM software should make use of commercial standards and NDI products to the maximum extent possible.

### **3.2.3.10 Instrumentation, Recording and Assessment**

This section provides rationale and guidance for instrumentation, recording and assessment services of OA components, processors, networks, etc. in support of component testing, certification and dynamic RM.

#### **3.2.3.10.1 Description**

Computer systems have long been instrumented, both during development and during deployment. The benefits provided to programmers and engineers for debug and testing during development are well established. The benefits during deployment are equally valuable in that recorded data may be used to diagnose anomalous situations experienced during operation.

With the advent of dynamic RM, the benefits of instrumentation take on an entirely new dimension; i.e., permitting real-time adaptation of the system configuration resulting from system health assessments based on instrumentation data. Two types of information are required for dynamic resource allocation: 1) application performance information and 2) computing resource consumption information. Resource consumption information is further subdivided into computer resource use and network resource use.

Application performance information is required to determining whether applications are meeting their weapon system performance requirements. Examples of this include periodic processing activation times and duration's, event latencies, and workload metrics (e.g., number of tracks processed). Status information is needed to determine the health of the application and whether the application is configured properly. Examples include initialization status, internal view of interface status, etc. Error condition information allows the reporting of exceptions and problems detected internally by the application. Examples include detection of corrupted data, faults on interfaces, middleware status or response problems, system call response errors, etc.

Computing resource consumption information is measured at the CPU level and at the network level. Resource use data provides the information needed to identify available resources that can be assigned when dynamically changing system configuration.

In fulfilling these instrumentation requirements, a set of instrumentation capabilities should be adopted that is consistent between the data collection needs and the RM needs.

#### **3.2.3.10.2 Guidance**

Guidance for instrumentation for OA is as follows:

1. Instrumentation should be used to collect run-time data and provide system performance feedback to RM monitoring capabilities. Services include data extraction, collection, reduction, analysis, and presentation.
2. The instrumentation capability should provide services and an API for collection of application-defined events and performance metrics. Information collected should include, but not be limited to, application performance and status data, application state and state changes, internal processing loads, occurrence of key events, internally detected error condition and internal view of interface status.
3. The instrumentation capability should provide services for measuring processor loading, memory use and other computing system metrics from all processors. Information collected should be at both the processor level and the process level. Where possible, thread information should also be collected. Information monitored should include:
  - 1) Operating system version and machine configuration
  - 2) CPU configuration, status, and utilization
  - 3) Memory configuration and usage
  - 4) Network configuration, status, and utilization
  - 5) File system configuration, status, and utilization
  - 6) Process status, including CPU, memory, network, and file system utilization for each process
  - 7) Network load as seen by a particular host
  - 8) Information concerning any remotely mounted file systems (e.g., NFS).
4. The instrumentation capability should provide services for measuring network metrics. Metrics should be collected at both the edge device connection level and on any internal switches, routers and their associated links. Metrics for each link should include, but are not limited to, bandwidth, throughput and latency. In addition, each host's view of its network loading should be collected (see item 3(7) above).
5. The instrumentation collection function should be implemented as a server that communicates with a scalable number of clients. Both the server and client processes should have minimal performance impact on the system under test.
6. Three levels of instrumentation should be defined: developmental, operational (in support of dynamic RM, see Section 3.2.3.8), and war diary. OA and resource allocation design should be such that operational and war diary instrumentation may be enabled at all times without adverse impact on tactical performance.
7. The server and the clients should not be developed with dependencies upon specific hardware items. The server should not limit the client's flexibility to manage the client's data buffer pools, control the events to monitor, specify data item formats, or group events into views of specific interest.
8. Details of services provided for monitoring and variations of collection software should be adapted to the available services where necessary.

### **3.2.3.11 Failure Management**

This section provides rationale and guidance for integrated failure management in the OACE (i.e., integration of fault detection, localization and recovery actions across all applications and supporting technologies).

#### **3.2.3.11.1 Description**

Failure management is the combined capabilities and interactions of all components of the architecture in support of fault tolerance for the overall system. It encompasses operating system services, middleware and associated system composition services, information transfer services, RM services, as well as application program design. This capability is particularly important since many component technologies present in the commercial marketplace are self-contained with respect to failure detection and fault tolerance.

#### **3.2.3.11.2 Guidance**

The following constitutes the failure management guidance for OACE components:

1. The failure management capability should operate in an integrated fashion across all of OACE components. The failure management capability should be designed in such a manner to ensure that detection and recovery actions should not cause erroneous system operation. This guidance should be relaxed in the case of network segmentation.
2. For COTS products that provide adequate access to internals, externally observable behavior should be monitored to the maximum extent practicable for the purpose of detecting and diagnosing faults and initiating recovery actions.
3. Failure management should consist of multiple tiers of recovery actions. The first tier should be built-in component failure recovery, whereby components initiate fail-over to replicated copies. The second tier should be re-instantiation of lost capability under control of RM. This would normally be accomplished by means of RM reload and restart of lost components. In some cases, where timeliness requirements permit, first tier failure recovery via application intrinsic replication may not be required.
4. OACE external interfaces to legacy components should utilize redundant physical connections between OACE communication nodes and external systems. The interface between the fully compliant OACE implementation and the associated combat system elements should provide for data flow and control behavior between the respective OACE subsystem and the rest of the weapon system.

#### **3.2.3.12 Information Assurance**

This section provides rationale and guidance for incorporation of information assurance features into OA and the OACE. It also addresses the interaction between information assurance technology and the resource manager.

##### **3.2.3.12.1 Description**

The overall objective of information assurance service is to ensure that mission critical applications have access to the necessary computing resources and data needed to complete the current mission of the ship successfully despite external attacks against the system. Traditional security services include authentication, confidentiality, and access control. Specifically, computationally expensive and complex encryption technologies that provide authentication and confidentiality services may not be needed inside the NWS enclave. These technologies may be needed on the “edges” of the enclave in order to support the encryption requirements of external

systems. The technology chosen to implement OA security measures should be readily certifiable.

The shift to CI/NDI components presents security threats that were not present in previous NWS baselines. The security vulnerabilities of the commercial world, such as denial of service attacks and computer viruses, are now a concern to the NWS since the common attacks found in the commercial world apply to many of the systems that will be used to host components of the NWS.

A defense in depth strategy should be considered in the design of OA. Firewalls, intrusion detection devices, and anti-virus software should be considered at points where the NWS communicates with external systems, whether these systems are shipboard or land-based.

#### **3.2.3.12.2 Guidance**

The following constitutes the guidance for information security interaction with RM:

1. The information assurance technology should have the ability to prevent known information warfare attacks.
2. The system should employ available information assurance technology to detect and counter information warfare attacks.
3. The information assurance technology should alert appropriate personnel and/or the resource manager when an information warfare attack has been detected.
4. The information assurance technology should notify, and if possible, recommend a response to appropriate personnel, security services and/or the resource manager when an information warfare attack has been detected.
5. The information assurance technology should be capable of implementing an automated response to an information warfare attack independent of personnel or the resource manager.

#### **3.2.3.13 Time Service**

This section provides rationale and guidance for providing a standard OA time distribution and synchronization service via the widely used commercial Network Time Protocol (NTP).

##### **3.2.3.13.1 Description**

The nature of real-time processing is that most, if not all, computers running real-time applications should be time synchronized to a fairly precise tolerance. This is true due to the coordination that should occur among the various weapon system components in the solution of problems involving physics and kinematics.

Two types of solutions are available to achieve time synchronization: 1) dedicated hardware (e.g., IRIG-B), and 2) software-based clock synchronization algorithms (e.g., NTP). Either approach is technically acceptable; however, the NTP approach has the advantage that it

does not require additional equipment. The NTP approach, which is inherently distributed and decentralized, is fairly robust.

Studies have shown that, with proper engineering of the time synchronization subnet, a one-millisecond time synchronization requirement can be met with commercial computers and networks using NTP. The core of NTP is a basic feedback loop controlled by a number of parameters including minimum and maximum polling interval. This feedback loop should be engineered on a system-by-system basis to achieve the required level of performance. This performance may be characterized by a number of factors including: 1) the level of synchronization required and 2) the time required to reach that level of synchronization within a defined standard deviation after a perturbation event (e.g., startup) has occurred.

This second parameter is sometime referred to as the NTP settling time and represents the time required for the feedback loop to stabilize after reset or loss of a clock. This settling time is dependent on a number of factors, including computer and operating system types, network characteristics and internal NTP parameter settings. Any decision concerning NTP use in the OA implementation should be based on engineering analysis of likely NTP settling time and synchronization accuracy compared to NWS operational requirements.

#### **3.2.3.13.2 Guidance**

Guidance for the OACE time service is as follows:

1. Support for a standardized combat-system-wide single time reference should be implemented in OA.
2. When performance characteristics permit, NTP should be considered for use as the standard time reference.
3. If a synthetic time service is used for embedded simulation and training operation, it should not interfere with or result in erroneous combat system operation. The artificial nature of this capability should be obvious to combat system operators and should be capable of rapid de-activation.

#### **3.2.3.14 Programming Language Facilities**

This section provides rationale and guidance for incorporation of programming language support facilities into the tool set for the OA development effort. The following programming languages are expected to be supported within OA applications:

- C/C++
- Java
- Ada

The guidance for programming language facilities is as follows:

1. Programming language facilities should support Procedural and Object Oriented Distributed programming.
2. All OACE infrastructure services should allow for bindings to all high-level languages used within the OA set of components.

3. Standard reuse libraries should be used where appropriate and effective.
4. Common coding and naming conventions should be applied throughout all OA software.

### **3.2.4 System Composition**

System composition is an aggregate term that describes synthesis of application components from available technologies, tools, models, methods, and mechanisms. It consists predominantly of utilizing support software, or “middleware,” to construct application computer programs in accord with design patterns such as those introduced in Section 3.2.1 and elaborated in Section 3.2.2.

Where middleware standards exist, products conforming to these standards should be used to construct application components. When standards do not yet exist, effort should be made during product selection to pick products that have widespread utilization and a commercial support structure.

For each application component, early design activities should focus on determining the functional capabilities and boundaries of each component. These activities include deciding what design pattern best fits the requirements, determining appropriate mechanisms for fault tolerance and scalability, identifying the interfaces with other components, and deciding on the appropriate technologies to utilize in effecting interactions between the component and other components.

### **3.2.5 Displays**

This section provides rationale and guidance for designing the OA display capability. It is recommended that the display architecture be designed so that any operator may use any console to accomplish any assigned mission.

#### **3.2.5.1 Description**

In legacy system implementations, display consoles were often unique to the operator they supported. In some instances, standardized consoles were customized through use of special purpose add-on devices and interfaces. This design approach was justifiable given the point-to-point interfaces and limitations in the graphics hardware and graphical support tools.

However, modern network and graphics technology provide an opportunity for a much more standardized and, therefore, logistically appealing solution in which any console may be used to perform the entire range of actions needed to support any operator mode or task. This approach not only provides logistics benefits, but it also promotes software reuse. Furthermore, the resultant equipment configuration is inherently more resilient in battle.

### **3.2.5.2 Guidance**

The following provides the guidance for display console organization with respect to generality of use.

1. Display and operator support software should conform to the guidance contained in this document.
2. Except for unique and/or special purpose requirements, the OA system should be designed such that each console supports the entire range of actions that need to be executed by each operator.

### **3.2.6 System Test and Certification**

This section provides rationale and guidelines for certifying systems based on the principles of dynamic resource allocation.

#### **3.2.6.1 Description**

Certification of traditionally designed systems is well understood. It consists primarily of functional and performance testing against specified system requirements, usually organized according to test cases and scenarios. However, systems employing dynamic management of resources, particularly those involving a degree of heterogeneity in the equipment configuration, require a more complex approach. Not only should the allocation mechanisms be tested, but also differences in equipment and configurations should be reconciled. This gives rise to additional requirements for certification of such systems.

In describing the guidance for certification of dynamically allocated systems, the term virtual homogeneity is used to describe a property of the computer and operating system that is required to support the certification process. This property states that each computer within a defined class of computers should perform substantially like every other computer in the class. Executables compiled for each computer in the class should yield operationally and statistically equivalent answers for an appropriate abstraction hierarchy (i.e. the granularity of system performance that is of interest to a particular user – in this case, the combat system's operators and sensor and weapons capabilities). Thus, virtual homogeneity is preserved when the variation in component output caused by any difference in computing resources is operationally negligible, a criterion that can be defined and tested.

Multiple classes of computers may be required for the foreseeable future to handle efficiently differing combat system requirements. For example, the type of computers and operating systems most effective for real-time weapons control may not be fully suitable for providing display capabilities and vice versa. However, several computers and operating systems may be fully capable of performing weapons control functions. The various computers falling in this latter category form an equivalence *class* and all products in this class are candidates for selection as OA computer resources.

This concept, when combined with the pool of computers concept discussed in Section 3.2.3.8 on RM, leads to multiple pools encompassing the equivalence classes needed to fully



implement NWS functionality. This constraint implies that several conditions should be met. First, a small number of computing resource types is preferable. Identical resources meet the virtual homogeneity criterion by definition, whereas a high degree of variation among processing resources will likely increase testing time and cost. In addition, sub-setting the total resource pool into classes of resources may result in a significant constraint on the dynamic RM.

If such constraints are needed, they could take the form of a tree of subsystems or aggregations of components. Each aggregation exists in only a manageable number of configurations. A case can be made that there are no (or negligible) interactions between the choices of configurations of sibling subsystems. Finally, all computers should provide a minimum level of performance (“unit of computing”) sufficient to ensure that system and component performance requirements are met.

Note that logistics considerations weigh in favor of limited numbers of classes and real homogeneity within classes.

Finally, any configuration that RM creates in the resource allocation process should be schedulable. Thus, part of the certification process should schedule issues. Also, requiring consideration during certification are stability considerations. In effect, RM’s allocation should not “trash” or produce a series of configuration changes sensitive to relatively small changes in the external environment. Due to the dynamic nature of RM, these considerations will likely require that support for the schedule and stability is incorporated directly into RM itself.

### **3.2.6.2 Guidance**

The guidance for system testing and certification of OA is as follows:

1. Traditional functional and performance testing should be done on have tactical and support components.
2. Functional and performance testing should be done on the dynamic RM capability.
3. Testing should be done to verify that system resources constitute “virtually homogeneous” pools.
4. If virtual homogeneity is not ensured across computing resource variants, additional testing is required to verify that resource allocation configuration constraints reflective of this are applied correctly.
5. Verification of scheduled analysis methods, stability analysis models, and capabilities should be built into the resource manager’s allocation methods.
6. All OACE components should have a regression test suite.

### **3.2.7 Selection of Standards**

A major advantage of the open system approach is that obsolete or failed components can be replaced quickly and easily. In order that this is accomplished effectively, it is highly advantageous to pick products in the commercial mainstream. Such products are far more likely to be readily available and interchangeable than niche market products.

The OACE should use industry standards to take maximum advantage of commercial mainstream technologies. Types of standards are listed in Table 5. Types of Standards below. Note that wherever possible, OACE-compliant system components should adhere to appropriate industry standards. There should be the minimum possible dependence of system components on vendor proprietary features. When features for which standards exist are employed, those standards should be employed. OACE components should be built and/or selected in a manner to foster portability across a variety of lower level technology base components. Metrics for measurement should be identified, and goal thresholds should be established for the metrics.

**Table 5. Types of Standards**

Standard Type	Definition
Formal Standard	Standards that are developed, approved, and maintained by a formally recognized standards committee by a consensus process. Examples: POSIX, CORBA.
Industry Standard	An industry standard is a formal or <i>de facto</i> standard that has been widely accepted and broadly implemented. A sufficient number of sources for products or services adhering to the standard are available to provide reasonable vendor independence. Examples: American National Standards Institute (ANSI) C, Small Computer System Interface (SCSI), PCI, and Java.
<i>Defacto</i> Standard	Standards that may have begun as proprietary but have achieved widespread acceptance by users. These standards are typically licensed to other vendors to produce compatible products. The vendor that produced the standard maintains the standard. Example: Parallel Computer Interface (PCI).
Proprietary Standard	Standards that have been published and are publicly available but the number of vendor implementations is limited (usually one). Example: Visual Basic. OACE should avoid the use of proprietary standards unless no other products will provide required functionality or performance.

### 3.2.7.1 Applicable Standards

A number of standards are sufficiently mature to merit their use in OA:

- IEEE POSIX P1003.1 2000 (currently in ballot, includes real-time and trace and P1003.2 shells and utilities)
  - P1003.1q Trace
  - P1003.1j Advanced Real-time Extensions
  - P1003.1d Additional Real-time Extensions
- P1003.25, Services for Reliable, Available, and Serviceable Systems (SRASS) (in ballot)
- IEEE P1003.21 Distributed Real-time services (in ballot)
- IEEE P1003.1 POSIX
- CORBA distributed object computing

- TCP/IP and UDP
- SQL
- XML

### **3.2.7.2 Open Source Products**

The open-source model provides an additional source of technology components for construction of the OACE. Open source products have the advantage of permitting the developer to review source code for quality control purposes and, if needed, add critical features vital to system operation. Some open source products do lack a commercial support base; however, many have a commercial support structure. Open source products should be considered when there are economic reasons for doing so and/or when such products provide a needed capability that is not currently obtainable by any other means. The guidance for open source products is as follows:

1. Wherever possible, open source products should have a commercial product support base.
2. Wherever possible, open source products should have a broad user base.
3. Licensing obligations for use of open source products should, not in any case, obligate the Navy or its contractors to return code to the open-source community; nor should use of open source product compromise the integrity of the Navy's combat system acquisition process.

### **3.2.7.3 Non-standard products**

Examples of non-standard but widely used products or products with unique capabilities that merit consideration – not all of which are real-time capable – include:

- VxWorks real-time operating system
- Microsoft Windows operating system
- Commercial database programs
- Commercial publish-subscribe protocols (e.g., NDDS)
- Ensemble group-ordered communication protocol
- Rosetta signal processing middleware
- DARPA Quorum ACE-operating system compatibility system
- DARPA Quorum/NSWCDD RM capability

(this page intentionally left blank)

# **APPENDIX A**

## **GLOSSARY OF TERMS**

AAW	Anti-Air Warfare
ACE	Adaptive Communications Environment
ACS	Aegis Combat System
ACTS	Aegis Combat Trainer System
ADS	Aegis Display System
ALIS	Aegis LAN Interconnect System
ANSI	American National Standards Institute
API	Application Program Interface
ARC	Aegis Reusable Components
ASN	Assistant Secretary of the Navy
ATDU	Aegis Tactical Display Upgrade
ATM	Asynchronous Transfer Mode
AWS	Aegis Weapon System
C&D	Command and Decision System
CC&D	Common Command and Decision
CCM	CORBA Component Model
CEC	Cooperative Engagement Capability
CG	Cruiser, Guided Missile
CI/NDI	Commercial Item / Non-Developmental Item
CNO	Chief Naval Operations
CORBA	Common Object Request Broker Architecture
COTS	Commercial-Off-The-Shelf
CSRD	Computing System Requirements Document
DARPA	Defense Advanced Research Projects Agency
DCOM	Distributed Component Object Model
DDG	Destroyer, Guided Missile
DII COE	Defense Information Infrastructure Common Operating Environment
DMS	Diminished Manufacturing Source
DoD	Department of Defense
DRM	Dynamic Resource Management
EDM	Engineering Development Model
FCS	Fire Control System
HiPer-D	Hi Performance Distributed Computing
HSI	Human Systems Integration
IDE	Integrated Development Environment
IDL	Interface Design Language
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Management Protocol
I/O	Input Output
IOP	I/O Processor
IPT	Integrated Product Team
IRIG-B	Inter-range Instrumentation Group

LAN	Local Area Network
LM/NE&SS	Lockheed Martin, Naval Electronic and Surveillance Systems
LTR	Letter
NAVSEA	Naval Sea Systems Command
NDDS	Network Data Distribution System
NFS	Network File System
NIC	Network Interface Card
NOFORN	No Foreign Releasability
NSWCDD	Naval Surface Warfare Center, Dahlgren Division
NTDS	Navy Tactical Data System
NTM	Network Transaction Management
NTP	Network Time Protocol
NTW	Navy Theater Wide
NWS	Naval Warfare Systems
OA	Open Architecture
OACE	Open Architecture Computing Environment
OMFTS	Operational Maneuver From the Sea
OMG	Object Management Group
ONR	Office of Naval Research
OO	Object Oriented
ORTS	Operational Readiness Test System
OS	Operating System
OSJTF	Open Systems Joint Task Force
PCI	Parallel Computer Interface
PEO-SC	Program Executive Office – Surface Combatants
PEO-TSC	Program Executive Office – Theater Surface Combatants
PIDS	Prime Item Development Specifications
PMS	Program Manager Sea (NAVSEA)
POSIX	Portable Operating System Interface (POSIX) Standard
QoS	Quality of Service
R&D	Research and Development
RD&A	Research, Development and Acquisition
RM	Resource Management
RMI	Remote Method Invocation
RSVP	Resource Reservation Protocol
SBC	Single Board Computer
SCIB	Ships Characteristics Improvement Board
SCSI	Small Computer System Interface
SECNAVINST	Secretary of the Navy Instruction
SER	Serial
SIG	Special Interest Group
SPAWARSYSCEN	Space and Naval Warfare Systems Center
SPY	Shipboard Phased Array Radar
SQL	Structured Query Language
SRASS	Services for Reliable, Available, and Serviceable Systems
SSDS	Ship Self Defense System

TAFIM	Technical Architecture for Information Management
TBMD	Theater Ballistic Missile Defense
TCP/IP	Transport Control Protocol / Internet Protocol
UDP	User Datagram Protocol
UDP/IP	User Datagram Protocol/Internet Protocol
UML	Unified Modeling Language
VME	Virtual Micro-bus European
WCS	Weapons Control System
XML	[E]Xtensible Markup Language